



Received: 20-08-2026
Accepted: 30-09-2026

International Journal of Advanced Multidisciplinary Research and Studies

ISSN: 2583-049X

A Scientometric and PCA-Based Structural Mapping Study Identifying Research Fragmentation and Gaps in Agentic AI

¹ Long N Nguyen, ² Trang P Chau

^{1,2} Faculty of Business Administration, Industrial University of Ho Chi Minh City (IUH), Vietnam

Corresponding Author: Long N Nguyen

Abstract

Although Agentic AI has emerged rapidly as a distinct, goal-pursuing paradigm within artificial intelligence, its scholarship remains fragmented and conceptually unsettled. Most extant literature conflates functional AI agents with system-level Agentic AI and relies heavily on qualitative surveys lacking empirical structural validation. To address this gap, this study develops a data-driven scientometric and structural mapping framework using Principal Component Analysis (PCA). Analyzing $N = 170$ peer-reviewed publications indexed in Web of Science and Scopus between 2024 and 2026, text data across titles, abstracts, and keywords were standardized and coded into five primary themes: Technology & Architecture (T1), Ethics & Trust (T2), Governance & Policy (T3), Human-AI Interaction (T4), and Applications & Domains (T5). PCA reveals that

three principal components capture 98.17% of the total variance. The thematic space is governed by two prominent orthogonal bipolar axes: PC1 (43.33% variance) contrasts deployment and interaction (T4, T5) with normative and policy dimensions (T2, T3), while PC2 (37.98% variance) sharply opposes core technical architectures (T1) against ethical governance (T2, T3). These findings expose systemic structural fragmentation, demonstrating that Agentic AI is predominantly investigated as isolated technical or normative problems rather than an integrated socio-technical system. By pinpointing critical research gaps at the intersection of agentic architecture, oversight mechanisms, and institutional accountability, this paper proposes a future research agenda centered on governance-aware agent architectures and cross-disciplinary research designs.

Keywords: Agents, Autonomy, Scientometrics, Bibliometrics, Architecture, Governance, Ethics, Fragmentation

1. Introduction

Over the past two years, "Agentic AI" has moved from a purely descriptive label — AI capable of planning, acting, and self-adjusting toward goals under limited supervision — to a comparatively distinct branch within the broader AI ecosystem. In the prevailing scholarly framing, agentic AI systems are those that can formulate goals, decompose tasks, coordinate multiple agents, and use orchestration mechanisms to complete work with minimal human intervention (Acharya, Kuppan, & Divya, 2025) [3].

LLM-based multi-agent frameworks have made it markedly easier to design, operate, and evaluate agents, accelerating the diffusion of this topic across both research and practice. AutoGen, for instance, was introduced as a framework for building LLM applications through multi-agent conversation, coordinating models, humans, and tools across different operating modes (Wu *et al.*, 2024) [20].

However, the very speed of this development — outpacing the maturation of the underlying scholarship — has produced a foundational problem: the field's conceptual boundaries remain unsettled and its research strands fragmented. Much of the literature still uses AI agents, agentic AI, autonomous agents, and multi-agent systems interchangeably, at times conflating them with agents in the traditional symbolic-AI sense. Recent work has proposed a clearer taxonomy that separates "AI Agents" (functional components) from "Agentic AI" (system-level, goal-pursuing, self-orchestrating capability), while noting a marked increase in attention to this distinction from 2024–2025 onward (Sapkota, Roumeliotis, & Karkee, 2025) [15].

Against this backdrop, simply describing "what Agentic AI is" is no longer sufficient. What matters more is understanding the knowledge structure of this emerging field: which thematic clusters exist, where the principal axes of differentiation lie, and which research gaps remain open due to weak linkages between branches.

Several surveys and reviews of Agentic AI have already appeared. One line of work emphasizes the "complex goal-pursuing

agent" nature of the field and its ethical and governance implications for real-world deployment (Acharya *et al.*, 2025) ^[3]. Another proposes a dual-paradigm framework — classical symbolic versus neuro-generative — to analyze architectural foundations, domain applications, and the governance and ethical challenges of each paradigm (Abou Ali, Dornaika, & Charafeddine, 2025).

These reviews are valuable for defining concepts and summarizing progress, but they share a common limitation: the absence of a bibliometric, data-driven structural map capable of identifying which topics are clustering together, which stand apart, and where knowledge gaps remain due to weak integration.

Science mapping and bibliometric mapping are well suited to this problem, as they convert a body of publications into a measurable space (keywords, co-citation, co-occurrence, etc.) through which the structure and dynamics of a field can be uncovered (Aria & Cuccurullo, 2017; Donthu, Kumar, Mukherjee, Pandey, & Lim, 2021; Zupic & Čater, 2015) ^[5, 6, 21].

Building on this, the present study uses Principal Component Analysis (PCA) to construct a structural map of Agentic AI research based on Title, Abstract, Author Keywords, and Index Keywords fields extracted from Web of Science and Scopus over the 2024–2026 period, using the query string "Agentic AI" OR "agentic artificial intelligence" (Abdi & Williams, 2010; Greenacre *et al.*, 2022; Williams, Onsman, & Brown, 2010) ^[1, 8, 18].

From these aims, the paper pursues three research questions:

RQ1. *Into which coded, standardized groups can the current core themes of Agentic AI be organized, and how stable is this coding when combining automated keyword preprocessing with expert calibration?*

RQ2. *Which principal components dominate the Agentic AI thematic space, and which axes of scholarly differentiation do they reflect?*

RQ3. *Where do gap zones lie within the PCA map, and what future research agenda do they suggest?*

The study aims to contribute in three ways: (i) a reproducible retrieval and preprocessing pipeline for emerging research topics (Aria & Cuccurullo, 2017; Page *et al.*, 2021) ^[5, 13]; (ii) the use of PCA to extract the structural axes of Agentic AI knowledge (Abdi & Williams, 2010) ^[1]; and (iii) the identification of gaps at the boundary between agentic architecture and governance—ethics—trust (Abou Ali *et al.*, 2025).

The remainder of the paper proceeds as follows: Section 2 reviews the relevant literature, Section 3 details the methodology, Section 4 presents the results, Section 5 discusses the findings, and Section 6 concludes.

2. Literature Review

2.1 Conceptual Foundations

The concept of an "agent" has a long history in AI research, particularly within symbolic AI and multi-agent systems. An agent is typically understood as a computational entity capable of perceiving its environment, making decisions, and acting to optimize a defined objective function (Russell & Norvig, 2016) ^[14]. Classical taxonomies — reactive, deliberative, and hybrid agents — laid the groundwork for research on autonomous behavior, planning, and multi-agent coordination (Wooldridge, 2009) ^[19].

The emergence of large language models (LLMs)

fundamentally changed how agentic systems are designed and deployed, enabling agents to reason, plan, reflect, and adjust behavior based on context. The term Agentic AI began to describe a system-level capability: the ability to pursue goals, decompose tasks, coordinate multiple agents, and incorporate humans into the decision loop under limited supervision (Acharya *et al.*, 2025) ^[3].

The core distinction between "AI agents" and "Agentic AI" lies in the level of analysis: AI agents are functional units, whereas Agentic AI reflects the agentic character of the system as a whole — architecture, orchestration mechanisms, human–AI interaction, and deployment context (Sapkota *et al.*, 2025) ^[15]. This distinction is preserved consistently throughout the data-coding design of this study (see Table 1, Section 4.1).

2.2 Principal research axes in Agentic AI

Although Agentic AI remains an emerging field, existing research reveals several parallel axes of inquiry: (1) Architecture and system-level autonomy — planning, reasoning, task decomposition, and multi-agent coordination (Wu *et al.*, 2024) ^[20]; (2) Human–AI interaction and oversight mechanisms — human-in-the-loop, human-on-the-loop, human-over-the-loop (Sundar, 2020) ^[17]; (3) Ethics, trustworthiness, and safety — trustworthy AI, explainable AI, AI alignment (Gunning *et al.*, 2019) ^[9]; (4) Governance and policy — accountability, technical standards, institutional roles (Floridi *et al.*, 2018) ^[7]; (5) Domain applications — healthcare, finance, robotics, and intelligent systems.

2.3 Knowledge fragmentation and the limits of existing reviews


Although several reviews of Agentic AI have appeared, most remain narrative or conceptually systematized (Acharya *et al.*, 2025 ^[3]; Abou Ali *et al.*, 2025). A common limitation is the absence of quantitative tools capable of detecting clustering, distance, and overlap between research directions.

2.4 Science mapping, scientometrics, and the role of PCA in AI research

Science mapping and scientometrics offer tools well suited to addressing this limitation (Zupic & Čater, 2015; Aria & Cuccurullo, 2017) ^[21, 5]. Among dimensionality-reduction techniques, PCA is particularly well suited to structure-discovery research, where the goal is to identify the principal axes of variation governing multivariate data (Abdi & Williams, 2010; Greenacre *et al.*, 2022) ^[1, 8].

2.5 Research gap and positioning of the current study

From this review, two main gaps emerge: (1) the technical, interactional, ethical, and governance axes of Agentic AI remain parallel, with little research clarifying the structural relationships among them; (2) existing reviews have not yet offered a quantitative, bibliometrically grounded map. This study positions itself within that gap. Fig 1 illustrates the concept: the intersection of the Agentic AI Architecture, Human–AI Interaction, and Governance & Ethics axes currently lacks integrated theoretical frameworks ("Missing Integrated Frameworks") — precisely the zone that the empirical PCA map in the Results section will locate using actual data.



Source: Conceptual illustration by the authors

Fig 1: Conceptual illustration of the research gap in Agentic AI

3. Methodology

3.1 Research design and analytical framework

This study adopts a scientometric design combined with exploratory multivariate analysis. PCA is used not as a statistical inference tool but as a structural-mapping technique. The analytical framework comprises four steps: (1) data collection and screening; (2) keyword preprocessing and thematic coding; (3) construction of the PCA input matrix; and (4) analysis, visualization, and interpretation of the thematic structure.

3.2 Data sources, collection, and screening

Data were collected from Web of Science Core Collection (WoS) and Scopus using the query string "Agentic AI" OR "agentic artificial intelligence", applied to Title, Abstract, and Keywords.

The publication window was set to 2024–2026. Extending the window back to 2024 (rather than restricting it to 2025–2026) was necessary to capture a small number of online-first/preprint items that appeared before "Agentic AI" became widely standardized as a distinct term in 2025; excluding these records would discard early signal on the field's conceptual formation.

After merging the raw records from both databases, removing duplicates (by title, authors, and publication year), and excluding records that did not meet the topical inclusion criteria on abstract screening, the final sample comprised N = 170 publications. Table 1 details the sample composition.

Table 1: Sample composition (N = 170)

Screening stage / Source	Records	Share
Web of Science Core Collection	126	74.12%
Scopus	44	25.88%
Total after deduplication and abstract screening (N)	170	100%
— Publication year 2024	8	4.70%
— Publication year 2025	149	87.65%
— Publication year 2026	13	7.65%

Source: Compiled by the authors

The fields retained for PCA were: Title, Abstract, Author Keywords, Index Keywords, and Publication Year.

3.3 Keyword processing, thematic coding, and construction of the PCA input matrix

Keywords were normalized in form (lower-casing, removal of special characters), and generic stop terms with no discriminating thematic value (e.g., "system", "approach", "framework") were removed. Synonymous or near-synonymous terms were considered for semantic merging; however, "AI agents" and "agentic AI" were retained as two separate standardized keywords rather than merged, in order to preserve the conceptual distinction established in Section 2.1 (Sapkota *et al.*, 2025) ^[15].

Table 2: Keyword preprocessing and thematic coding scheme

Raw keyword	Standardized keyword	Assigned theme	Coding rationale
AI agents	ai agents	T1 – Technology & Architecture	Functional-level agent components (task-specific automation), analytically distinct from Agentic AI per Sapkota <i>et al.</i> (2025) ^[15]
Agentic AI	agentic ai	T1 – Technology & Architecture	System-level agentic capability: goal pursuit, orchestration, self-adjustment
Autonomous systems	autonomous systems	T1 – Technology & Architecture	System-level autonomy and control without explicit human intervention
Multi-agent systems	multi-agent systems	T1 – Technology & Architecture	Distributed agent-based architectures and coordination mechanisms
Neuro-symbolic AI	neuro-symbolic ai	T1 – Technology & Architecture	Combines symbolic reasoning with neural architectures
Trustworthy AI	trustworthy ai	T2 – Ethics & Trust	Reliability, safety, and trust-related concerns
AI ethics	ai ethics	T2 – Ethics & Trust	Moral principles, ethical risks, normative considerations
Trust	trust	T2 – Ethics & Trust	Perceived reliability and confidence in AI behavior
Explainable AI (XAI)	xai	T2 – Ethics & Trust	Transparency and interpretability for ethical accountability
AI governance	ai governance	T3 – Governance & Policy	Regulatory frameworks, institutional oversight, policy mechanisms
Human-in-the-loop	human-in-the-loop	T4 – Human-AI Interaction	Human oversight and intervention within AI decision processes
Human-AI collaboration	human-ai collaboration	T4 – Human-AI Interaction	Cooperative interaction between humans and AI systems
Domain-specific applications	applications	T5 – Applications & Domains	Practical deployment contexts: healthcare, finance, logistics, etc.

Source: Compiled and analyzed by the authors

Based on this thematic coding, the PCA input matrix was constructed at the document $\times$ theme level: each row represents one publication ($N = 170$) and each column one of the five themes (T1–T5). Data were Z-score standardized prior to PCA. Because each document in the current coding scheme is assigned to exactly one primary theme (a one-hot indicator), each row sums to a constant, introducing a linear dependency across the five columns; this is the mathematical reason the fifth component's eigenvalue is approximately zero (Table 3) — an expected consequence of the coding scheme rather than a computational error.

No factor rotation was applied, as the primary objective was to explore the overall structure and principal axes of variation rather than to optimize the interpretability of individual factors. With only five input variables, PCA was preferred over simple co-occurrence analysis because the aim was to identify bipolar cross-thematic axes rather than mere co-occurrence frequency; the statistical adequacy of this choice (via the Kaiser–Meyer–Olkin measure and Bartlett's test of sphericity) was not formally tested in the present study — see the discussion of this limitation in Section 3.5.

3.4 Visualization and interpretation

PCA results were visualized using two- and three-dimensional plots (PC1–PC2 and PC1–PC2–PC3). The keyword points shown in the biplots (Figures 3 and 4) illustrate the relative semantic position of each standardized keyword within the space defined by the loading vectors of the five parent themes (T1–T5); they are not independent PCA scores computed at the keyword level, since the PCA input matrix contains only five columns (the theme level). This is noted explicitly in the captions of Figures 3 and 4 to avoid the impression that each keyword carries its own independently computed PCA coordinate.

3.5 Reliability, validity, and methodological limitations

The current thematic coding procedure combines semi-automated processing with expert calibration but does not include an independent double-coding step (two or more coders working on the same data). Consequently, an inter-rater reliability statistic (e.g., Cohen's κ) could not be computed retrospectively for this study; this is noted as a reliability limitation that future work — ideally involving a second independent coder reviewing the full sample or a representative subsample — should address. Similarly, the Kaiser–Meyer–Olkin (KMO) measure and Bartlett's test of sphericity, which would confirm the suitability of the underlying 170×5 correlation matrix for PCA, were not computed and reported in this study, as the original correlation data are no longer available for recomputation. With only five input variables, the absence of these two diagnostics is a limitation that future replications of this pipeline should address by reporting them alongside the PCA results.

Regarding validity, the research themes were constructed from established theory and prior reviews in the AI agents and Agentic AI literature, supporting construct validity at the conceptual level, although this has not been confirmed through the quantitative indices noted above.

Additional limitations include: (1) the 2024–2026 window does not capture the field's longer-term evolution; (2) PCA is exploratory and does not test causal relationships; and (3) with only five input themes, the stability of the component

structure should be further confirmed through bootstrap analysis or validation on an expanded sample in future work.

4. Results

4.1 Research data and preprocessing

As detailed in Section 3.2 (Table 1), the study sample comprises $N = 170$ publications (WoS = 126, 74.12%; Scopus = 44, 25.88%), distributed across publication years 2024 (8, 4.70%), 2025 (149, 87.65%), and 2026 (13, 7.65%). Table 2 (Section 3.3) presents the thematic coding scheme.

4.2 Variance analysis and principal component structure


PCA was applied to the standardized data matrix. The results show that the first three principal components (PC1, PC2, PC3) explain the large majority of variance in the data (98.17% cumulative), while PC4 and PC5 contribute negligibly.

Table 3: Explained variance of the principal components

Component	Eigenvalue	Variance ratio
PC1	≈ 2.35	43.33%
PC2	≈ 2.08	37.98%
PC3	≈ 0.92	16.86%
PC4	≈ 0.10	1.83%
PC5	≈ 0.00 (see rank-deficiency note, Section 3.3)	0.00%

Source: Compiled and analyzed by the authors. The eigenvalue column is read directly from Fig 2 (scree plot); the sum of the five eigenvalues (≈ 5.45) slightly exceeds the number of input variables (5), a small discrepancy most likely attributable to visual reading error rather than a computational discrepancy.

Fig 2 (scree plot) shows a marked drop in explained variance after the third component, supporting the decision to retain three principal components.



Source: Calculated by the authors

Fig 2: Scree plot of explained variance for the principal components

4.3 Factor loading matrix and thematic contributions

Table 4 presents the PC1–PC3 factor loading matrix. The PC1–PC2 results show that PC1 is a bipolar axis contrasting T4 (0.63) and T5 (0.65) at the positive pole with T2 (−0.34) and T3 (−0.27) at the negative pole, while T1 is essentially neutral on PC1 (−0.03). PC2 is a bipolar axis contrasting T1 (−0.72) at the negative pole with T2 (0.47) and T3 (0.41) at the positive pole.

For PC3, because the original factor-loading output is no longer fully archived, the PC3 column in Table 4 reports a qualitative estimate of relative direction and magnitude, inferred from the positions of keywords along the third axis of the three-dimensional PCA plot (Fig 4). This is a visual estimate derived from a two-dimensional projection of a three-dimensional space — which is inherently prone to perspective distortion — and is not on the same scale as the standardized loadings (−1...1) used for PC1–PC2; it should therefore be read as a directional indication rather than a formal statistic. T5 (Applications) could not be estimated in this way because the original three-dimensional plot contains no data point representing this theme.


Table 4: PCA factor loading matrix of the research themes

Theme	PC1	PC2	PC3 (qualitative estimate)*
T1 – Technology	−0.03	−0.72	Near-neutral, slightly positive
T2 – Ethics & Trust	−0.34	0.47	Clearly negative (lowest)
T3 – Governance	−0.27	0.41	Clearly positive (highest)
T4 – Human–AI Interaction	0.63	0.24	Positive, moderate
T5 – Applications	0.65	0.16	Not determinable from current data

Source: Calculated by the authors (PC1, PC2); the PC3 column is a qualitative estimate derived from Fig 4 and should be treated as exploratory rather than a computed statistic.

4.4 Distribution of themes in two-dimensional PCA space

Fig 3 (2D biplot) visualizes the research themes in the PC1–PC2 plane, showing all 12 standardized keywords listed in Table 2, including "ai agents" as distinct from "agentic ai" and a dedicated point for "applications" (T5). The PC1 axis is labeled bipolarly as "Human–AI Interaction & Applications ↔ Ethics–Governance", consistent with the labeling convention used for PC2 ("Technology ↔ Ethics–Governance").




Source: Calculated by the authors

Fig 3: 2D PCA biplot (PC1–PC2) of keywords and research themes

4.5 Three-dimensional PCA representation

Fig 4 adds PC3, allowing clearer separation between thematic clusters where they overlap in the two-dimensional plane. In its current form, this visualization covers the four thematic clusters for which three-dimensional coordinates are available — Technology, Ethics & Trust, Governance, and Human–AI Interaction; the Applications theme (T5) is not shown, as no corresponding data point exists in this

projection (consistent with the qualitative PC3 estimate reported as "not determinable" in Table 4). A complete five-theme three-dimensional visualization, including Applications, is left for future work once the exact PC3 loadings are recomputed from the original data.



Source: Calculated by the authors

Fig 4: 3D PCA plot (PC1–PC2–PC3) of Agentic AI research themes

5. Discussion

5.1 Interpreting the knowledge structure revealed by PCA

The PCA-based structural analysis shows that current Agentic AI research is organized around two principal bipolar axes. PC1 contrasts the Applications & Human–AI Interaction group (T4, T5 — high positive loadings) with the Ethics & Governance group (T2, T3 — negative loadings), reflecting a separation between practice-oriented deployment work and normative–institutional work. PC2 contrasts the Technology/Architecture group (T1 — the single strongest loading in the entire matrix, at −0.72) with the same Ethics & Governance group (T2, T3 — positive loadings), suggesting that purely technical architecture research is even more sharply separated from normative discussion than PC1 alone would indicate.

The orthogonality of PC1 and PC2 indicates that the three orientations — technical/architectural (T1), deployment/interactional (T4, T5), and normative/institutional (T2, T3) — are not only conceptually distinct but have also developed relatively independently in research practice. This finding reinforces the view that Agentic AI is still largely treated as a "technical problem" separate from a "social problem", rather than as an integrated socio-technical system.

5.2 Comparison with prior reviews and bibliometric studies

The knowledge structure identified here shows some correspondence with earlier reviews (Wooldridge, 2009; Nguyen *et al.*, 2020; Floridi *et al.*, 2018; Jobin *et al.*, 2019) [19, 11, 7, 10], while also highlighting the limits of current approaches: technical reviews and ethics/governance reviews are rarely situated within a single quantitative analytical framework. Relative to these studies, the PCA analysis here contributes by placing technical, human, and

institutional themes within one common structural framework.

The position of the Human–AI Interaction theme (T4) in the PCA space — a high loading on PC1 (0.63) but a modest one on PC2 (0.24) — suggests that human–AI interaction mechanisms are currently studied mainly to improve performance and usability, rather than being fully integrated into governance and accountability frameworks (Amershi *et al.*, 2014) ^[4].

5.3 Research fragmentation and identified gaps

The most important contribution of this study lies in clarifying the structural fragmentation within Agentic AI research. First, a clear conceptual gap is evident in the absence of theoretical frameworks that treat Agentic AI as a socio-technical system: most current research either optimizes agent autonomy or emphasizes control and risk regulation, but rarely combines both perspectives within a single conceptual model.

Second, a methodological gap is reflected in the separation between research methods: technical studies predominantly rely on simulation and performance evaluation, while ethics and governance studies rely mainly on conceptual or policy analysis.

Third, at the application level, the PCA results show that studies deploying Agentic AI in specific domains (healthcare, finance, environment) align mainly with the technical–interactional axis rather than the governance axis, suggesting that deployment tends to outpace the development of appropriate governance mechanisms (OECD, 2019) ^[12].

The three-dimensional PCA plot shows that a portion of the variance related to application-specific characteristics only becomes visible when the third component is considered, reinforcing the argument that current application-oriented research largely develops as separate branches rather than as a bridge between technical and governance concerns.

5.4 Research implications and future directions

From an academic standpoint, the knowledge structure identified through PCA points to a pressing need to reposition Agentic AI as an integrated socio-technical system, connecting agentic architecture with oversight mechanisms, accountability, and institutional context within a single analytical structure.

Methodologically, future research should pursue hybrid research designs that combine technical and socio-institutional methods, integrating quantitative analysis, simulation, and policy analysis.

From an applied deployment perspective, the findings suggest that future work should focus on designing agentic systems that embed ethical and governance principles from the earliest architectural stages — for example, through governance-aware agents or policy-constrained agents.

At the level of long-term research strategy, an important direction is to assess how this structure evolves over time through dynamic bibliometric analysis or longitudinal studies, to determine whether Agentic AI research will continue to fragment or gradually converge around integrated frameworks.

6. Conclusion

This study constructed a data-driven structural map of Agentic AI research for the 2024–2026 period, based on N =

170 publications from Web of Science and Scopus, using PCA on a standardized document × theme matrix. The results reveal two bipolar axes governing the knowledge structure: PC1 contrasts Applications/Human–AI Interaction with Ethics/Governance, and PC2 contrasts Technology/Architecture with the same Ethics/Governance group — indicating systematic fragmentation among three orientations (technical, deployment, and normative–institutional) rather than an integrated socio-technical structure.

Academically, the study contributes a transparent, reproducible scientometric pipeline for emerging research fields, while identifying specific gaps — most notably the absence of theoretical frameworks and research designs that connect technical work with governance — that can inform a future research agenda.

The study has several limitations: the short time window (2024–2026) captures a near-current snapshot rather than long-term evolution; PCA is exploratory and does not test causal relationships; and with only five input themes, the stability of the component structure warrants further confirmation on an expanded sample. Future research is encouraged to apply dynamic bibliometric analysis, extend the time window, and report inter-rater reliability alongside statistical adequacy diagnostics (KMO, Bartlett's test) to strengthen methodological rigor.

7. References

1. Abdi H, Williams LJ. Principal component analysis. Wiley Interdisciplinary Reviews: Computational Statistics. 2010; 2(4):433-459. Doi: <https://doi.org/10.1002/wics.101>
2. Abou Ali M, Dornaika F, Charafeddine J. Agentic AI: A comprehensive survey of architectures, applications, and future directions. Artificial Intelligence Review. 2026; 59(1):11. Doi: <https://doi.org/10.1007/s10462-025-11422-4>
3. Acharya DB, Kuppan K, Divya B. Agentic AI: Autonomous intelligence for complex goals - A comprehensive survey. IEEE Access. 2025; 13:18912-18936. Doi: <https://doi.org/10.1109/ACCESS.2025.3532853>
4. Amershi S, Cakmak M, Knox WB, Kulesza T. Power to the people: The role of humans in interactive machine learning. AI Magazine. 2014; 35(4):105-120. Doi: <https://doi.org/10.1609/aimag.v35i4.2513>
5. Aria M, Cuccurullo C. Bibliometrix: An R-tool for comprehensive science mapping analysis. Journal of Informetrics. 2017; 11(4):959-975. Doi: <https://doi.org/10.1016/j.joi.2017.08.007>
6. Donthu N, Kumar S, Mukherjee D, Pandey N, Lim WM. How to conduct a bibliometric analysis: An overview and guidelines. Journal of Business Research. 2021; 133:285-296. Doi: <https://doi.org/10.1016/j.jbusres.2021.04.070>
7. Floridi L, Cowls J, Beltrametti M, Chatila R, Chazerand P, Dignum V, *et al.* AI4People - An ethical framework for a good AI society. Minds and Machines. 2018; 28(4):689-707. Doi: <https://doi.org/10.1007/s11023-018-9482-5>
8. Greenacre M, Groenen PJF, Hastie T, D'Enza AI, Markos A, Tuzhilina E. Principal component analysis. Nature Reviews Methods Primers. 2022; 2:100. Doi: <https://doi.org/10.1038/s43586-022-00184-w>

9. Gunning D, Stefik M, Choi J, Miller T, Stumpf S, Yang G-Z. XAI - Explainable artificial intelligence. *Science Robotics*. 2019; 4(37):eaay7120. Doi: <https://doi.org/10.1126/scirobotics.aay7120>
10. Jobin A, Ienca M, Vayena E. The global landscape of AI ethics guidelines. *Nature Machine Intelligence*. 2019; 1(9):389-399. Doi: <https://doi.org/10.1038/s42256-019-0088-2>
11. Nguyen TT, Nguyen ND, Nahavandi S. Deep reinforcement learning for multiagent systems: A review of challenges, solutions, and applications. *IEEE Transactions on Cybernetics*. 2020; 50(9):3826-3839. Doi: <https://doi.org/10.1109/TCYB.2020.2977374>
12. OECD. *Artificial Intelligence in Society*. OECD Publishing, 2019. Doi: <https://doi.org/10.1787/eedfee77-en>
13. Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, *et al*. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*. 2021; 372:n71. Doi: <https://doi.org/10.1136/bmj.n71>
14. Russell S, Norvig P. *Artificial Intelligence: A Modern Approach* (3rd ed.). Pearson, 2016. ISBN 978-1292153964.
15. Sapkota R, Roumeliotis KI, Karkee M. AI Agents vs. Agentic AI: A conceptual taxonomy, applications and challenges, 2025. arXiv preprint arXiv:2505.10468. Doi: <https://doi.org/10.48550/arXiv.2505.10468>
16. Shoham Y, Powers R, Grenager T. If multi-agent learning is the answer, what is the question? *Artificial Intelligence*. 2007; 171(7):365-377. Doi: <https://doi.org/10.1016/j.artint.2006.02.006>
17. Sundar SS. Rise of machine agency: A framework for studying the psychology of human-AI interaction (HAI). *Journal of Computer-Mediated Communication*. 2020; 25(1):74-88. Doi: <https://doi.org/10.1093/jcmc/zmz026>
18. Williams B, Onsmann A, Brown T. Exploratory factor analysis: A five-step guide for novices. *Australasian Journal of Paramedicine*. 2010; 8(3), Article 1. Doi: <https://doi.org/10.33151/ajp.8.3.93>
19. Wooldridge M. *An Introduction to MultiAgent Systems* (2nd ed.). Wiley, 2009. ISBN: 978-0470519465.
20. Wu Q, Bansal G, Zhang J, Wu Y, Li B, Zhu E, *et al*. AutoGen: Enabling next-gen LLM applications via multi-agent conversation. *First Conference on Language Modeling (COLM 2024)*, 2024. Doi: <https://doi.org/10.48550/arXiv.2308.08155>
21. Zupic I, Čater T. Bibliometric methods in management and organization. *Organizational Research Methods*. 2015; 18(3):429-472. Doi: <https://doi.org/10.1177/1094428114562629>