

Mapping the Research Landscape of Agentic Artificial Intelligence: A Bibliometric Analysis of Scopus Publications

Zafar Imam Khan

Learning Resources Manager, Department of Library, Hamdan Bin Mohammed Smart University, Dubai, United Arab Emirates

Email: zafarimamkhan@gmail.com | ORCID: <https://orcid.org/0000-0003-2081-0951>

Received: 31 August 2026 | Revised: 19 September 2026 | Revised: 21 September 2026 | Accepted: 26 September 2026 | Published online: 30 September 2026

How to cite: Khan, Z. I. (2026). Mapping the research landscape of agentic artificial intelligence: A bibliometric analysis of Scopus publications. *Journal of Scholarly Research and Metrics*, 1(1), 5–26.

Abstract

Context: Agentic Artificial Intelligence (Agentic AI) is rapidly emerging as a major research paradigm in which artificial intelligence systems move beyond response generation toward autonomous, goal-directed reasoning, planning, tool use, memory, adaptation, and multi-agent collaboration. The rapid expansion and disciplinary fragmentation of this literature create a need for systematic mapping of its research landscape.

Aim: This study maps the global development, scholarly influence, geographical distribution, and collaboration patterns of Agentic AI research during 2023–2026.

Approach: A bibliometric and science-mapping analysis was conducted using Scopus publications retrieved on 1 July 2026. Following systematic screening and duplicate removal, 3,425 publications were analyzed using Biblioshiny.

Key Findings: Agentic AI exhibited an annual growth rate of 215.04%, with 9,166 authors contributing across 1,015 sources. Conference papers dominated the corpus (2,654), reflecting the field's rapidly evolving technological character. China and the United States were the leading corresponding-author countries, while Tsinghua University was the most productive affiliation. Citation analysis identified several highly influential foundational studies.

Interpretation: The findings indicate that Agentic AI is developing rapidly into a distinct, collaborative, and increasingly established research domain centered on autonomous and coordinated AI systems.

Contribution: This study provides a broad Scopus-based quantitative map of contemporary Agentic AI scholarship and identifies emerging priorities for evaluation, multi-agent coordination, security, explainability, human–AI collaboration, and governance.

Keywords: agentic AI; AI agents; large language models; bibliometric analysis; science mapping; Scopus

1. Introduction

AI systems are shifting from forecasting or reacting toward acting autonomously to achieve specified goals. This new paradigm, called Agentic Artificial Intelligence (Agentic AI), encompasses planning, reasoning, memory, tool use, interaction with the environment, adaptation,

and autonomous decision-making with limited continuous human supervision (Acharya et al., 2025; Pati, 2025). This change has been fueled by recent advancements in large language models (LLMs), which offer a cognitive backbone for agents to reason, plan, communicate, and engage with evolving environments (Stübinger & Metz, 2026).

While standalone language models are capable of generating content, LLM-driven agents can incorporate components such as memory, planning capabilities, access to external tools, and feedback loops to enhance their performance. Their key abilities include perception, decision-making, action, and iterative learning (Li & Zhu, 2025). Thus, Agentic AI is being explored in a wide variety of application fields, such as software engineering, healthcare, robotics, education, communication systems, and more (Karunanayake, 2025; Otoum & Elkhaili, 2026).

Scholarly interest has also expanded rapidly. Bibliometric evidence already identifies several major and emerging research directions in LLM-based intelligent agents, including architecture and technical foundations; multi-agent collaboration; system optimization; domain applications; safety; trustworthiness; and ethical governance (Li & Zhu, 2025). However, the literature remains conceptually and disciplinarily fragmented. Previous bibliometric studies have often focused on specific applications, such as multi-agent formation control (Xu et al., 2025b), autonomous robotics (Casini et al., 2025), tourism (Le, 2026), and small, medium, and micro enterprises (SMMEs) (Olujimi et al., 2025).

Accordingly, a broader quantitative mapping of contemporary Agentic AI research is needed. The present study analyzes 3,425 Scopus-indexed publications from 2023 to 2026, examining scientific production, influential publication sources and documents, institutional productivity, geographical contributions, international collaboration, and citation impact.

1.1 Research Question and Overview of the Study

In light of the growth and fragmentation of Agentic AI scholarship, this study poses the following broad research question:

RQ: How has the global research landscape of Agentic Artificial Intelligence evolved from January 2023 to 1 July 2026 in terms of scientific production, leading sources and contributors, geographical distribution, collaboration patterns, and citation impact?

To answer this question, 3,425 publications were analyzed by combining performance analysis and science-mapping techniques. The findings indicate a rapidly growing field, with the number of documents published increasing from 41 in 2023 to 1,696 in 2025, while the 2026 count reflects only publications indexed up to the data-collection date. China and the United States are the dominant contributing countries, and the most productive institutions are concentrated mainly in China and other technologically advanced research settings. The reviewed literature also points to a transition from basic LLM studies toward more autonomous, cooperative, and integrated agentic systems. These patterns are examined in detail in the following sections.

2. Literature Review

2.1 Evolution and Conceptual Foundations of Agentic Artificial Intelligence

The conceptual foundations of Agentic Artificial Intelligence (Agentic AI) are based on the development of AI from rule-based, task-specific computational systems to increasingly autonomous, adaptive and goal-oriented systems. Early AI focused mainly on narrowly defined tasks and was subject to explicit human control; behavior was defined with explicit rules and representations. Unlike traditional AI, which typically responds to each stimulus individually, modern AI agents are capable of adapting their strategies and taking actions to achieve goals within a changing world (Moralles et al., 2026). This advancement is not merely a boost in model

performance but a transformation in how AI is perceived, from a computational instrument to a system that can display aspects of agency.

The evolution of agency in AI can be understood through several technological and theoretical stages. An important early foundation was symbolic artificial intelligence and expert systems, in which knowledge representations, inference mechanisms and explicit rules were used to mimic parts of the human problem-solving process. These systems, however, were usually reactive and lacked the ability to proactively modify their behavior due to changes in the environment. Further research on autonomous agents, specifically Belief–Desire–Intention (BDI) architectures, proposed a more explicit notion of agency, modeling agents as having beliefs or knowledge, desired goals, and commitments, commonly referred to as plans or intentions. These advances led to the emergence of artificial agents that, unlike traditional expert systems, are able to create goals, plan actions and modify their behavior based on feedback (Moralles et al., 2026).

Building on this, reinforcement learning, evolutionary computation, and multi-agent systems further extended this trajectory, enhancing the ability of artificial agents to learn from their environments, adapt to uncertainty, and collaborate with other computational agents. Multi-agent systems were especially relevant because intelligence was no longer limited to a single decision-making entity but extended to communication, cooperation, competition, and the division of duties among multiple entities. Similarly, Sapkota et al. (2026) identify earlier research on expert systems and multi-agent systems as key predecessors of current AI agents. They note that autonomy, environmental perception, communication, distributed problem solving, and social interaction are key antecedents of contemporary agentic systems. Therefore, Agentic AI should not be regarded as a completely new technological creation, but as a new stage in the development of long-standing ideas: autonomous agents, cognitive architectures, distributed intelligence, planning, and adaptive learning.

However, modern-day Agentic AI is very different from many traditional autonomous-agent architectures. A useful distinction can be made between symbolic/classical agency and neural/generative agency. Abou Ali et al. (2025) regard these as two separate but increasingly linked lineages. The symbolic/classical lineage is based mainly on explicit representation, algorithmic planning, stateful memory, and static reasoning mechanisms, while the neural/generative lineage is based on stochastic generation, foundation models, prompt-based reasoning, and dynamic orchestration. The distinction is valuable because it addresses how agency is produced, which can be difficult to capture using traditional symbolic-agent terminology. In modern neural systems, agent-like behavior emerges through prompting, contextual interaction, tool invocation, memory, and orchestration rather than through explicitly encoded symbolic beliefs and intentions (Abou Ali et al., 2025).

The emergence of generative AI and LLMs has led to a significant shift. Generative models offer flexible natural-language understanding, knowledge synthesis, and reasoning capabilities. Agentic systems extend this paradigm by enabling a model to pursue more complex, multi-stage goals. Bandi et al. (2025) define Agentic AI as a new frontier that amalgamates the capabilities of autonomy, adaptability, planning, memory, reflection, and goal-driven reasoning, going beyond conventional AI and generative AI. An agentic system can not only answer a prompt but also interpret what it is supposed to do, set sub-goals, select an appropriate set of tools, take a series of actions, monitor intermediate results, and proceed largely autonomously.

Therefore, autonomy is not enough to define Agentic AI. According to Acharya et al. (2025), the paradigm is characterized by autonomous systems aimed at complex goals, which need to be performed with limited human presence and show adaptation to changing environments, advanced decision-making ability, and self-sufficiency. Their conception also marks a shift from passive, reactive systems to AI systems that can plan strategically, process information and solve problems even under conditions that are not fully determinable beforehand. These qualities indicate that agency should be considered as a multi-faceted construct of autonomy, goal

orientation, environmental responsiveness, temporal continuity, reasoning, planning, memory, and adaptive action rather than a single technical property.

AI agents are generally associated with modular systems that leverage models and external tools for task-specific automation, while Agentic AI is more comprehensive, featuring dynamic task decomposition, memory, coordinated autonomy, and cooperation between multiple agents. Abou Ali et al. (2025) provide a complementary perspective: an AI agent can be considered as an autonomous unit that can achieve a goal, whereas Agentic AI is the broader principle of structuring agency, often via teams of specialized agents communicating and coordinating with each other to work on tasks that go beyond the capabilities of any single agent. This distinction is especially relevant because the field has shifted from single-agent execution to system-level orchestration.

Architecturally, modern Agentic AI brings together a variety of interrelated features and abilities, not a single algorithm or model. These include the perception of inputs and environmental states, reasoning, memory, task decomposition, tool use, action execution, feedback processing, reflection and inter-agent coordination. Bandi et al. (2025) highlight that agentic systems can combine foundation models with planners, knowledge bases, sensors and actuators, external tools and feedback mechanisms to accomplish their goals while adjusting their plans in real time. Similarly, in modern multi-agent architectures, intelligence is increasingly distributed across agents, with orchestration mechanisms coordinating them, marking a move from a single model toward an ensemble of behaviors (Sapkota et al., 2026).

Despite increasing convergence around these capabilities, the concept of Agentic AI remains incompletely defined. Moralles et al. (2026) note that there is presently no universal definition and that terms like autonomous agents, intelligent agents and Agentic AI are used interchangeably or in contradictory ways. Two views of Agentic AI can be distinguished: one sees Agentic AI as a logical extension of the autonomous agents, while another sees it as an emergent property arising from the combination of LLMs with memory, feedback loops, external tools, and execution mechanisms. On this basis, a convenient way of conceptualizing Agentic AI for the purposes of this study is to see it as a general paradigm of AI in which computational systems have enough autonomy and organization to understand instructions, generate and/or break down tasks, reason and plan, converse with tools (or environments), remember and use context, consider and act on feedback, modify future steps, and, when necessary, coordinate with other agents while continuously interacting with them. This synthesis captures the elements most consistently discussed in the literature, while recognizing that the boundaries of the field continue to evolve.

2.2 Large Language Models and the Emergence of Agentic AI

The advent of Large Language Models (LLMs) has greatly spurred the shift from traditional autonomous agents to modern Agentic AI. Most earlier autonomous agents were built for limited environments, relied on predefined policies or environment-specific training, or had limited knowledge, restricting their ability to perform open-ended tasks. In contrast, LLMs offer general pre-trained knowledge together with natural-language understanding and reasoning, enabling them to serve as decision-making components in autonomous systems (Wang et al., 2024). This change has shifted the focus of language models from text generation to planning, decision-making, and goal-oriented action.

One of the key elements of this shift is the incorporation of LLMs into agent architectures instead of standalone models. Wang et al. (2024) contend that autonomous behavior requires more than question-answering capability: profile, memory, planning, and action modules are also needed. Memory enables agents to learn from past interactions, planning enables them to organize future actions, and action mechanisms convert the model's decisions into interactions with the external environment. Therefore, Agentic AI can be seen as a stepping stone from language generation to reasoning–planning–acting integration.

This process has been further strengthened by reasoning and task decomposition. Li (2025) describes methods that allow LLM-driven agents to break down complex problems and consider

multiple reasoning paths, including Chain-of-Thoughts (CoT), Tree-of-Thoughts (ToT), and Graph-of-Thoughts (GoT). LLMs are also increasingly connected to symbolic planners, external knowledge, APIs, retrieval mechanisms, and executable tools so that agents have the capability to go beyond the knowledge in their internal models to execute real tasks in the outside world (Li, 2025). These are essential components of Agentic AI because the goal of autonomy is not merely to generate a response but to decide which actions to take, in what order, and with what resources.

The importance of LLMs is also evident in embodied environments. Raptis et al. (2025) demonstrate that LLMs can be integrated into robotic systems to improve their high-level reasoning capabilities, task decomposition, planning, decision making, and adaptability. In such systems, LLMs can act as high-level cognitive controllers, with the rest of the robotic system performing low-level physical actions, allowing for the integration of language-based reasoning with environmental perception and action.

However, Agentic AI involves more than adding autonomy to an LLM. Haidemariam (2026) suggests that the agentic transition is an ongoing cycle of perception, evaluation, goal revision, and action, which brings AI closer to reflective and goal-driven reasoning. Effective Agentic AI is, therefore, enabled by orchestrating LLMs with memory, planning, tools, feedback, environmental interaction and adaptive decision-making mechanisms. This convergence marks the key shift from language models that produce information to agentic models capable of reasoning about goals and acting toward them.

2.3 Autonomous and Multi-Agent AI Systems

Autonomous AI systems extend large language models (LLMs) from output generation to goal-driven agents capable of reasoning, planning, memory, tool use, environmental interaction, and learning from feedback. Wang et al. (2024) introduce the idea of LLM-based autonomous agents with a central LLM and interconnected components: profiling, memory, planning, and action components. Recent reviews also focus on the perception, reasoning, planning, memory, execution, communication, and integration of external tools as prerequisites for autonomous operation (Alva & Pandey, 2026; Chowa et al., 2026; Ferrag et al., 2026). Systematic evaluation is also growing, as illustrated by Chowa et al.'s (2026) review of 68 publicly available datasets for evaluating LLM agents.

A major step in distributed intelligence is the move from individual autonomous agents to LLM-based multi-agent systems (MAS). In MAS, rather than a single general-purpose agent handling the whole problem, the problem is divided into multiple subtasks and these subtasks are handled by multiple specialized agents that communicate, coordinate, negotiate or jointly reason to achieve common goals (Brohi et al., 2025; Li et al., 2024). Li et al. (2024) outline the five stages of the LLM-based MAS: profile, perception, self-action, mutual interaction, and evolution, and highlight two major application domains: problem solving and world simulation. This supports better modularization and task distribution and, in turn, specialized architectures. For instance, SALLMA follows a two-layered architecture: an Operational Layer for orchestration in real-time, and a Knowledge Layer to store configurations of the workflows and the agents; the proof of concept uses Docker, Kubernetes, Python, LangChain, Hugging Face, Mistral, LLaMA, and SQL/NoSQL databases (Becattini et al., 2025).

Coordination structure is especially important, as more agents do not necessarily produce more effective results. MultiAgentBench is a suite of multi-agent collaboration and competition benchmarks that are based on six interactive scenarios and four different communication topologies: star, chain, tree and graph. Its experiments showed that GPT-4o-mini had the best average task score; graph-based coordination had the best performance in the research scenario; and cognitive planning had a 3% higher milestone achievement rate (Zhu et al., 2025). Domain-specific evidence also exists for robotics-related applications (Raptis et al., 2025), chemistry and automated experimentation (Ramos et al., 2025), software engineering (Becattini et al., 2025), and autonomous decision environments (Ferrag et al., 2026). For example, Ramos et al. (2025) report

the GenQEX benchmark of 457 questions, with over 20 genes identified by the evaluated TAIS case with > 80% prediction accuracy.

However, multi-agent autonomy also carries risks, including communication overhead, inter-agent misalignment, cascading errors, scalability, security, and unpredictable emergent behavior (Brohi et al., 2025; Chowa et al., 2026; Ferrag et al., 2026). Thus, while more capable LLMs are crucial to the development of autonomous and multi-agent AI, so are appropriate agent specialization, orchestration, and communication, together with effective tool integration, performance evaluation, and governance.

2.4 Applications and Emerging Research Directions in Agentic AI

Agentic AI is making rapid progress from experimental architectures to domain-specific systems with autonomous reasoning, planning, tool use and multi-step execution. Applications are emerging in fields such as healthcare, finance, cybersecurity, robotics, autonomous transportation, education, enterprise automation, smart cities, software engineering, and knowledge management (Brohi et al., 2025; Deshmukh et al., 2025; Pati, 2025). In each of these areas, the key change is that, beyond offering AI-driven recommendations, the new systems can coordinate actions and adapt workflows with less human oversight.

One of the most well-established application fields is healthcare. Research is examining agentic systems for electronic health record (EHR) interaction, clinical triage, medical question answering, diagnosis, patient monitoring, mental health support, and medical reasoning (Gridach et al., 2026). However, most healthcare agentic systems remain research prototypes: Gridach et al. (2026) found that none of the surveyed frameworks had been formally approved, prospectively validated in clinical settings, or incorporated into production EHRs, highlighting the need for deployment-focused research, formal safety controls, and regulatory planning.

Enterprise applications include meeting summarization, supply-chain optimization, customer-service orchestration, and employee-productivity workflows (Xu et al., 2025a); autonomous sales agents support lead generation, lead qualification, customer interaction, and performance management (Gonzalez et al., 2026). Broader assessments highlight other fast-growing deployment areas, such as digital twins, education, e-commerce, cyber security and network management (Khalid et al., 2026).

Emerging research is therefore focusing more on reliability, scalability, evaluation, and governance than on agent capability alone. In a study of 39 open-source agent frameworks and 439 agentic applications, Hasan et al. (2026) showed that over 70% of agent testing focuses on deterministic tools and workflows, while less than 5% addresses foundation-model-based planning and only about 1% tests prompts. In response, Akshathala et al. (2026) suggest evaluating beyond task completion under four headings: LLM, memory, tools, and environment. Hybrid neuro-symbolic agents, multimodal and embodied intelligence, lifelong learning, scalable multi-agent coordination, human-AI collaboration, security, explainability, and formal governance are among the other areas gaining importance for future research (Abou Ali et al., 2025; Hosseini & Seilani, 2025).

2.5 Bibliometric Research and Knowledge Gaps in Agentic AI

Evidence on Agentic AI is growing rapidly but remains fragmented across databases, application domains, and conceptual definitions. Existing research has focused on specific aspects of the field and has not yet produced a coherent picture of its overall knowledge structure. For instance, Zhang and Zhang (2026) evaluated 810 Web of Science publications from 2023 to 2025, which included 733 articles and 77 reviews, and found that the number of publications had increased from 4 in 2023 to 710 in 2025. In contrast, a bibliometric investigation of LLM-powered intelligent agents analyzed 620 publications, revealing three major research hot spots: agent architectures/technical foundations, multi-agent collaboration/optimization, and domain-specific applications (Li & Zhu, 2025).

Other bibliometric studies have focused on individual sectors. Olujimi et al. (2025) used VOSviewer and the Bibliometrix R package to analyze 214 papers on agentic AI in SMMEs published from 2019 to 2025, 162 of which appeared between 2023 and mid-2025. Their analysis showed concentrations of research in the United States, China, and the United Kingdom, as well as gaps in empirical implementation, governance, regional adaptation, and interdisciplinary research. In contrast, Badiu et al. (2025) used co-occurrence and co-citation clustering in VOSviewer to map research on AI agents and business performance across 106 articles published between 2019 and 2024.

The same specialization is evident in education. Zrigui et al. (2026) conducted a systematic bibliometric review of generative AI agents in personalized learning and assessment, retaining 1,206 publications and 38 empirical studies for in-depth review, and highlighted reproducibility concerns such as insufficient reporting of agent architectures and system configurations. Ugur (2026) similarly applied bibliometric analysis to agentic AI in education and proposed an agentic pedagogical agency framework.

Other bibliometric research likewise reveals the disciplinary dispersion of this literature. Xu et al. (2025b) studied 1,851 Web of Science publications about multi-agent formation control and found that four themes—cooperation, formation tracking, obstacle avoidance, and system stability—were of particular interest. Casini et al. (2025) used SciMAT and VOSviewer for Scopus literature mining related to human-centered AI and autonomous robotics with a focus on reliability, safety, trustworthiness and human–robot collaboration. While these studies offer valuable insights, their narrow disciplinary focus does not capture the full scope of the Agentic AI research landscape.

3. Research Methodology

3.1 Research Design

This study employed a quantitative bibliometric and science-mapping design to examine the development, performance, and structure of Agentic AI research between 1 January 2023 and 1 July 2026. To capture the current wave of LLM-powered Agentic AI, the analysis begins in 2023. Although autonomous agents, intelligent agents, and multi-agent systems have a longer research history, the advent of large language models from 2023 onward significantly advanced research on agents that combine LLM-powered reasoning, planning, memory, tool use, and autonomous action. The study therefore focuses on this contemporary strand of research rather than taking a broader historical view of autonomous and multi-agent systems. Performance analysis was combined with collaboration and citation analyses to provide a multidimensional view of the research landscape. Because data were retrieved on 1 July 2026, the 2026 records represent partial-year coverage and are interpreted accordingly.

3.2 Data Source and Search Strategy

Bibliographic data were retrieved exclusively from Scopus on 1 July 2026. The search was conducted within titles, abstracts, and keywords (TITLE-ABS-KEY) using terms representing Agentic AI and LLM-enabled agents, including “agentic AI,” “AI agent*,” “LLM agent*,” “autonomous agent*,” “generative agent*,” “multi-agent LLM*,” and “agentic system*,” combined with “large language model*,” LLM*, “generative AI,” ChatGPT, GPT-4, and GPT-3.5. The initial search identified 6,851 records.

Final search string: TITLE-ABS-KEY (("agentic AI" OR "AI agent*" OR "LLM agent*" OR "LLM-based agent*" OR "LLM-powered agent*" OR "language agent*" OR "autonomous agent*" OR "generative agent*" OR "multi-agent LLM*" OR "agentic workflow*" OR "agentic system*") AND ("large language model*" OR LLM* OR "generative AI" OR ChatGPT OR GPT-4 OR GPT-3.5)) AND (EXCLUDE (DOCTYPE , "cr") OR EXCLUDE (DOCTYPE , "ch") OR EXCLUDE (DOCTYPE , "no") OR EXCLUDE (DOCTYPE , "bk") OR EXCLUDE (

DOCTYPE , "ed") OR EXCLUDE (DOCTYPE , "le") OR EXCLUDE (DOCTYPE , "sh") OR EXCLUDE (DOCTYPE , "er") OR EXCLUDE (DOCTYPE , "dp")) AND (LIMIT-TO (LANGUAGE , "English")) AND (LIMIT-TO (PUBSTAGE , "final")) AND (EXCLUDE (SRCTYPE , "k") OR EXCLUDE (SRCTYPE , "d")) AND (EXCLUDE (SUBJAREA , "VETE") OR EXCLUDE (SUBJAREA , "IMMU") OR EXCLUDE (SUBJAREA , "PHAR") OR EXCLUDE (SUBJAREA , "NURS") OR EXCLUDE (SUBJAREA , "AGRI") OR EXCLUDE (SUBJAREA , "ECON") OR EXCLUDE (SUBJAREA , "HEAL") OR EXCLUDE (SUBJAREA , "MULT") OR EXCLUDE (SUBJAREA , "NEUR") OR EXCLUDE (SUBJAREA , "PSYC") OR EXCLUDE (SUBJAREA , "EART") OR EXCLUDE (SUBJAREA , "CHEM") OR EXCLUDE (SUBJAREA , "CENG") OR EXCLUDE (SUBJAREA , "ENVI") OR EXCLUDE (SUBJAREA , "BIOC") OR EXCLUDE (SUBJAREA , "BUSI") OR EXCLUDE (SUBJAREA , "ENER") OR EXCLUDE (SUBJAREA , "MATE") OR EXCLUDE (SUBJAREA , "PHYS") OR EXCLUDE (SUBJAREA , "MEDI") OR EXCLUDE (SUBJAREA , "SOCI") OR EXCLUDE (SUBJAREA , "ARTS")) AND PUBYEAR > 2022 AND PUBYEAR < 2027

Scopus was selected as the sole bibliographic database because of its broad multidisciplinary coverage and structured bibliographic metadata suitable for performance analysis and science mapping, including publication, citation, authorship, affiliation, and keyword information. The search was restricted to English-language records to maintain consistency in the analysis and interpretation of bibliographic and keyword metadata. The disciplinary scope was limited to Computer Science, Engineering, Mathematics, and Decision Sciences because the study primarily focuses on the computational and technological development of contemporary Agentic AI and LLM-enabled agent systems. Although Agentic AI has applications across numerous disciplines, these four subject areas were selected to maintain a focused corpus aligned with the study's technological research objectives. The potential exclusion of relevant research indexed in other disciplinary categories is acknowledged as a limitation of the study.

3.3 Eligibility and Screening

Table 1 presents the PRISMA-based study identification and screening process employed in this research. The initial Scopus search conducted on 1 July 2026 retrieved 6,851 records. Sequential eligibility criteria were then applied for publication year, document type, language, publication stage, source type, and subject area, as specified in Table 1. The analysis retained publications from January 2023 to 1 July 2026, comprising conference papers, articles, and reviews published in English with final publication status and indexed within the selected subject areas. Metadata-quality screening and duplicate removal were subsequently performed before construction of the final bibliometric dataset.

Table 1. PRISMA-Based Study Identification and Screening Process

Step	Screening Criterion	Included	Excluded	Remaining Records
1	Initial Scopus search (1 July 2026) using the predefined search string	6,851	–	6,851
2	Publication year (2023–2026)	2023–2026	2022 (2), 2027 (5)	6,844
3	Document type	Conference Paper,	Conference Review,	6,432

		Article, Review	Book Chapter, Note, Book, Editorial, Letter, Short Survey, Erratum, Data Paper (412)	
4	Language	English	Non-English (81)	6,351
5	Publication stage	Final	Article in Press (190)	6,161
6	Source type	Conference Proceedings, Journals	Book Series (513), Trade Journals (5)	5,643
7	Subject area	Computer Science, Engineering, Mathematics, Decision Sciences	All remaining subject areas (2,170)	3,473
8	Metadata quality and title-based relevance screening	Records with usable essential metadata and titles relevant to the study scope	Records with missing/unusable essential metadata, structural metadata problems, or clearly irrelevant titles (38)	3,435
9	Duplicate removal	Unique records	Duplicate records (10)	3,425

Source: Author's compilation based on Scopus search and screening data.

3.4 Data Cleaning and Final Dataset

After database-level filtering, 3,473 records were exported and examined in Microsoft Excel for metadata completeness, relevance, structural consistency, and duplication. The completeness of metadata was evaluated based on essential bibliographic elements needed for the next step of analysis. Records with missing or unusable information were considered incomplete. Records with obviously inconsistent or incomplete bibliographic data were also eliminated if the data could not be safely used for bibliometric analysis.

Relevance screening was performed mostly on the title of the publications. If a record was retrieved by the search strategy but its title was clearly not indicative of it being focused on contemporary Agentic AI and LLM-based agent research, it was considered irrelevant and removed. The purpose of this study was not to systematically review individual publications but to carry out a bibliometric mapping; therefore, full-text screening was not performed. As this is a single-author study, the metadata-quality and title-based relevance screening was performed by the author. These procedures removed 38 records, leaving 3,435.

The remaining records were then checked for duplicates, and 10 duplicate records were removed. The final analytic dataset thus consisted of 2,654 conference papers, 718 articles and 53 reviews, for a total of 3,425 unique publications. All performance-analysis and science-mapping procedures were executed based on this cleaned dataset.

3.5 Bibliometric Analysis

Annual scientific production, the sources of publications, Bradford's law, affiliations, countries of corresponding authors, international collaboration, impact of citations, and global citations of documents were analyzed. The final corpus comprised 1,015 sources and 9,166 authors, growing at a rate of 215.04% each year.

3.6 Analytical Software

The cleaned dataset was subsequently analyzed using Biblioshiny, the web-based interface of the Bibliometrix package implemented in R. Biblioshiny/Bibliometrix was used to generate descriptive bibliometric indicators and conduct analyses of scientific production, Bradford's Law, institutional productivity, geographical contribution, collaboration, and citation impact. The analytical environment used for the final analysis and verification was R version 4.6.1 with Bibliometrix version 5.5.0.

3.7 Visualization and Figure Preparation

All quantitative results, numerical values, and bibliometric relationships presented in the figures were derived directly from the Scopus dataset and Biblioshiny/Bibliometrix analyses. The original Biblioshiny outputs were retained as the reference source for figure preparation. ChatGPT was used only as an AI-assisted tool to improve the graphical presentation of selected outputs, including readability, font visibility, spacing, layout, and publication-quality formatting. ChatGPT was not used to generate, recalculate, infer, or modify any bibliometric values, rankings, relationships, or analytical findings. Following graphical enhancement, each figure was manually cross-checked against the corresponding original Biblioshiny output to verify the accuracy of numerical values, labels, rankings, and represented relationships. Thus, AI assistance was limited to visual presentation and did not form part of the bibliometric analysis or interpretation.

4. Results

Table 2. Descriptive characteristics of the Agentic AI bibliometric dataset

Indicator	Value
Dataset profile	
Study period	2023–2026
Publications analyzed	3,425
Publication sources	1,015
Annual growth rate	215.04%
Mean document age	0.77 years
Mean citations per publication	5.76
Total references	216,581
Knowledge descriptors	
Keywords Plus	14,113
Author keywords	6,825
Authorship profile	
Total authors	9,166
Authors of single-authored publications	200
Single-authored publications	214
Mean co-authors per publication	4.8
International co-authorship	27.30%
Publication types	
Conference papers	2,654

Articles	718
Reviews	53


Note: The final corpus of 3,425 publications is consistent with the dataset obtained after metadata screening and duplicate removal in the study-selection procedure.

Source: Author's analysis based on Scopus data using Biblioshiny (Bibliometrix R package).

Table 2 portrays Agentic AI as a rapidly expanding and strongly collaborative research domain. Covering 2023–2026, the corpus comprises 3,425 publications distributed across 1,015 sources, with an annual growth rate of 215.04%, indicating substantial acceleration in scholarly activity during the study period. The relatively low mean document age of 0.77 years further reflects the recency of this research landscape. Despite its emerging nature, the literature has already accumulated an average of 5.76 citations per publication and draws on 216,581 references, suggesting both early scholarly influence and extensive engagement with existing knowledge.

The intellectual diversity of the field is reflected by 6,825 author keywords and 14,113 Keywords Plus terms. Research production is also highly collaborative: 9,166 authors contributed to the corpus, with an average of 4.80 co-authors per publication, while 27.30% of publications involved international co-authorship. Only 214 documents were single-authored, reinforcing the collaborative character of Agentic AI research. Regarding publication format, conference papers were predominant (2,654), considerably exceeding articles (718) and reviews (53). This conference-oriented publication pattern is consistent with a fast-developing technological research area in which emerging methods, architectures, and applications are communicated rapidly through conference venues.

Figure 1. Annual scientific production, 2023–1 July 2026



Source: Author's analysis based on Scopus data using Biblioshiny (Bibliometrix R package).

The annual scientific production shown in Figure 1 demonstrates rapid growth in Agentic AI research during the study period. Publication output increased substantially from 2023 through 2025, indicating accelerating scholarly attention to LLM-enabled agents, autonomous systems, and related Agentic AI technologies. The dataset contains 1,282 publications indexed for 2026 as of the data-retrieval date of 1 July 2026. This value represents partial-year coverage and should therefore not be interpreted as evidence of a decline in annual research output relative to 2025. Additional publications may be indexed during the remainder of 2026, and indexing delays may further affect the observed count. Accordingly, comparisons involving 2026 are treated cautiously throughout the study.


Table 3. Core Publication Sources According to Bradford’s Law

Rank	Core Source	Publications	Cumulative Publications	Bradford Zone
1	Proceedings of the AAAI Conference on Artificial Intelligence	144	144	Zone 1
2	AAMAS 2026 – Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems	138	282	Zone 1
3	Conference on Human Factors in Computing Systems – Proceedings	112	394	Zone 1
4	CEUR Workshop Proceedings	98	492	Zone 1
5	Proceedings of Machine Learning Research	83	575	Zone 1
6	Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)	62	637	Zone 1
7	Advances in Neural Information Processing Systems	56	693	Zone 1
8	International Conference on Agents and Artificial Intelligence	43	736	Zone 1
9	Proceedings of the IEEE International Conference on Big Data	38	774	Zone 1
10	NAACL-HLT 2025: Long Papers	37	811	Zone 1

Source: Author’s analysis based on Scopus data using Biblioshiny (Bibliometrix R package).

Table 3 identifies the core publication sources according to Bradford’s Law. The Proceedings of the AAAI Conference on Artificial Intelligence ranked first with 144 publications, followed closely by AAMAS 2026 with 138 and the CHI proceedings with 112. Collectively, the ten Zone 1 sources contributed 811 publications, highlighting a concentrated core of influential publication venues. The strong representation of conference proceedings further indicates the conference-driven nature of knowledge dissemination in Agentic AI research.


Figure 2. Most Productive Affiliations in Agentic AI Research



Source: Author's analysis based on Scopus data using Biblioshiny (Bibliometrix R package).

Figure 2 illustrates the ten most productive institutional affiliations contributing to Agentic AI research during 2023–2026. Tsinghua University leads with 102 publications, followed by Zhejiang University (86) and Nanyang Technological University (69). The University of California contributed 60 publications, while Shanghai Jiao Tong University recorded 58. Fudan University, Carnegie Mellon University, Peking University, the National University of Singapore, and the University of Chinese Academy of Sciences each contributed between 47 and 53 publications. The distribution highlights the prominent role of Asian institutions, particularly Chinese universities, in shaping the emerging Agentic AI research landscape.


Figure 3. Leading Corresponding Author Countries and International Collaboration in Agentic AI Research



Source: Author's analysis based on Scopus data using Biblioshiny (Bibliometrix R package).

Figure 3 presents the leading countries based on corresponding-author publications and their collaboration patterns. China dominates with 771 publications (22.5%), followed by the USA with 513 (15.0%). Although their publication volumes are lower, Singapore (59.6%), Canada (45.3%), and the United Kingdom (43.2%) exhibit the highest proportions of multiple-country publications (MCP), indicating strong international collaboration. In contrast, Japan records the lowest MCP rate (12.7%) among the leading countries. Overall, the findings reveal substantial geographical concentration in China and the USA alongside varying levels of cross-national research collaboration.

Figure 4. Citation Impact of Leading Countries in Agentic AI Research



Source: Author's analysis based on Scopus data using Biblioshiny (Bibliometrix R package).

Figure 4 depicts the citation performance of the ten most cited countries in Agentic AI research. China achieved the highest overall citation impact with 5,304 total citations, followed by the USA with 2,438 citations. Singapore and Switzerland recorded 668 and 658 citations, respectively. In terms of average citations per article, Iran exhibited the highest value (153.00), followed by Switzerland (34.60) and Singapore (11.70). The findings demonstrate a clear distinction between overall citation volume and average citation impact, with China and the USA dominating total citations while several lower-output countries exhibit comparatively high citation intensity.

Table 4. Most Globally Cited Documents in Agentic AI Research

Rank	Document	DOI	Total Citations	Citations/Year	Normalized TC
1	Park JS (2023), UIST – Proc. ACM Symp. User Interface Softw. Technol.	10.1145/3586183.3606763	1,766	441.5	14.42
2	Shinn N (2023), Adv. Neural Inf. Process. Syst.	—	1,381	345.25	11.28
3	Wang L (2024), Front. Comput. Sci.	10.1007/s11704-024-40231-1	1,365	455	60.4
4	Bran AM (2024), Nat. Mach. Intell.	10.1038/s42256-024-00832-8	531	177	23.49
5	Shen Y (2023), Adv. Neural Inf. Process. Syst.	—	507	126.75	4.14
6	Guo T (2024), IJCAI Int. Joint Conf. Artif. Intell.	—	255	85	11.28
7	Zhang R (2024), IEEE J. Sel. Areas Commun.	10.1109/JSAC.2024.3459037	218	72.67	9.65
8	Ge Y (2023), Adv. Neural Inf. Process. Syst.	—	205	51.25	1.67
9	Zhao A (2024), Proc. AAAI Conf. Artif. Intell.	10.1609/aaai.v38i17.29936	191	63.67	8.45
10	Mandi Z (2024), Proc. IEEE Int. Conf. Robot. Autom.	10.1109/ICRA57147.2024.10610855	165	55	7.3

Source: Author's analysis based on Scopus data using Biblioshiny (Bibliometrix R package).

Table 4 highlights the most globally cited documents within the analyzed Agentic AI literature. Park et al. (2023) leads with 1,766 citations, followed by Shinn et al. (2023) with 1,381 and Wang et al. (2024) with 1,365. Notably, Wang et al. records the highest citation rate (455 citations per year) and normalized citation score (60.40), indicating particularly strong impact relative to publications from the same period. Overall, the concentration of highly cited works from 2023–2024 demonstrates the rapid scholarly uptake of foundational research during the early development of Agentic AI.

5. Discussion

This study provides a broad quantitative assessment of the rapidly growing landscape of Agentic Artificial Intelligence research. Analysis of 3,425 Scopus-indexed publications from 2023 to 2026

demonstrates substantial publication growth, a high level of scholarly collaboration, conference-oriented knowledge dissemination, geographical concentration, and rapidly developing citation impact. The annual growth rate of 215.04%, together with contributions from 9,166 authors across 1,015 publication sources, reflects the rapid expansion of Agentic AI as a contemporary research domain.

5.1 Rapid Growth and Maturation of Agentic AI Research

This growth is well demonstrated by the annual publications. Scientific production increased from 41 publications in 2023 to 406 in 2024 and 1,696 in 2025. The 1,282 publications recorded for 2026 do not indicate a decline in scientific activity, because data were retrieved only up to 1 July 2026 and 2026 is therefore a partial publication year. This rapid development aligns with recent work describing Agentic AI as a paradigm shift from traditional generative models toward systems that combine autonomy, planning, memory, reasoning, tool use, and goal-directed action (Acharya et al., 2025; Bandi et al., 2025; Moralles et al., 2026).

This rapid expansion coincides with the growing scholarly interest in LLM-based agent architectures. Recent literature describes LLMs as increasingly being incorporated into systems that combine reasoning with memory, planning, external tools, feedback, and environmental interaction (Wang et al., 2024).

5.2 The Structure of Publications and Scholarly Communication

The publication profile is dominated by conference papers (2,654), followed by journal articles (718) and reviews (53). This pattern is confirmed by the Bradford's Law analysis, which shows that the 10 core sources account for 811 publications, with the Proceedings of the AAAI Conference on Artificial Intelligence as the top source with 144 publications, followed by AAMAS 2026 with 138 publications and CHI proceedings with 112 publications. Other important conferences are NeurIPS, PMLR, CEUR Workshop Proceedings and IEEE Big Data.

This conference-oriented structure aligns with the technology-driven nature of Agentic AI, in which new architectures, benchmarks, frameworks, and applications emerge rapidly. It also suggests that the field has not yet shifted toward the journal-based publication patterns typical of more established fields. The strong representation of AAAI, AAMAS, NeurIPS, and CHI in the publication corpus also reflects the interdisciplinary character of Agentic AI research, spanning artificial intelligence, autonomous and multi-agent systems, machine learning, and human-computer interaction.

5.3 Geographical Leadership and Research Collaboration

The institutional and geographical results show a high concentration of Agentic AI research. The three most productive affiliations were Tsinghua University with 102 publications, Zhejiang University with 86, and Nanyang Technological University with 69. Other Chinese institutions that featured in the top ten most productive affiliations were Shanghai Jiao Tong University, Fudan University, Peking University and the University of Chinese Academy of Sciences.

At the country level, China (771) and the United States (513) produced substantially more corresponding-author publications than other leading countries. However, patterns of international collaboration differed from publication volume: the share of multiple-country publications was highest in Singapore (59.6%), followed by Canada (45.3%) and the United Kingdom (43.2%). Agentic AI research thus combines strong domestic research capacity in leading countries with an active international collaboration network. This trend follows previous bibliometric research that points to the United States, China, and the United Kingdom as key geographical sources of more specific contributions in Agentic AI areas (Olujimi et al., 2025).

5.4 Citation Impact and Intellectual Foundations

Citation analysis identifies several highly influential contributions in the early Agentic AI literature. Park et al. (2023) accumulated the highest total citation count (1,766), followed by Shinn et al. (2023) with 1,381 and Wang et al. (2024) with 1,365. Importantly, Wang et al. (2024) has the highest citation rate (455 citations per year) and normalized citation score (60.40), indicating particularly strong influence on a per-year basis.

Country-level citation patterns also reveal differences between research volume and citation intensity. China, the most productive country, received substantially more citations (5,304) than the United States (2,438), the second most productive country, while some countries with relatively few publications showed very high average citation rates. These averages should be interpreted with caution, as they may be skewed by a small number of highly cited publications.

5.5 Implications and Future Research Directions

Overall, the results suggest a turning point in Agentic AI research: the key question is no longer whether LLM-based agents can demonstrate agent-like behavior, but how agentic systems can be reliably designed, coordinated, evaluated, secured, and governed. This interpretation aligns with the recent literature on the key research areas of reliability, scalability, multimodal intelligence, embodied intelligence, human–AI collaboration, explainability, security, evaluation, and governance (Abou Ali et al., 2025; Hosseini & Seilani, 2025; Khalid et al., 2026).

The study also addresses a significant gap in the bibliometric literature. Previous bibliometric studies have often focused on specific applications, such as tourism (Le, 2026), autonomous robotics (Casini et al., 2025), small, medium, and micro enterprises (SMMEs) (Olujimi et al., 2025), and multi-agent formation control (Xu et al., 2025b). By examining 3,425 publications across scientific production, publication sources, institutional and geographical contributions, collaboration patterns, and citation impact, the present study provides a broader quantitative perspective on contemporary Agentic AI research than previous domain-specific bibliometric investigations. Future research should further explore the critical issues of robust agent evaluation, multi-agent coordination, human oversight, security and failure propagation, explainability, governance, and longitudinal changes in the geographical and institutional landscape, as well as the possibility that the current geographical and institutional concentration will spread across the globe as Agentic AI evolves.

6. Conclusion

This study provides a comprehensive bibliometric assessment of 3,425 Scopus-indexed publications on Agentic Artificial Intelligence from January 2023 to 1 July 2026. The findings demonstrate rapid growth in scholarly activity, extensive collaboration, and a predominantly conference-driven publication landscape. China and the United States emerged as the leading corresponding-author countries, while several Asian institutions occupied prominent positions in research productivity. Citation analysis identified a small group of foundational studies that have rapidly achieved substantial scholarly influence. Overall, the findings demonstrate the rapid development and increasing scholarly visibility of Agentic AI as a contemporary research domain. Future research should further examine robust agent evaluation, scalable multi-agent coordination, security, explainability, human–AI collaboration, governance, and changes in geographical and institutional participation as the field develops.

Declarations

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflict of Interest

The author declares no conflict of interest.

Author Contributions (CRediT)

Zafar Imam Khan: Conceptualization, Methodology, Data Curation, Formal Analysis, Visualization, Investigation, Writing – Original Draft, Writing – Review & Editing.

Ethics Approval and Informed Consent

Not applicable. This study is based exclusively on bibliographic data obtained from Scopus and did not involve human participants or animals.

Declaration of Generative AI Use

ChatGPT (OpenAI) was used solely to improve the visual presentation of selected figures (readability, font visibility, spacing, and layout), as described in Section 3.7. It was not used to generate, recalculate, or modify any bibliometric data, results, or interpretations. The author verified all figures against the original Biblioshiny outputs and takes full responsibility for the content of this article.

Data Availability

The bibliographic data analyzed in this study were retrieved from the Scopus database on 1 July 2026. The analyzed data may be made available by the author upon reasonable request, subject to Scopus licensing and database-access conditions.

Acknowledgements

Not applicable.

References

- Abou Ali, M., Dornaika, F., & Charafeddine, J. (2025). Agentic AI: A comprehensive survey of architectures, applications, and future directions. *Artificial Intelligence Review*, 59(1), 11. <https://doi.org/10.1007/s10462-025-11422-4>
- Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic AI: Autonomous intelligence for complex goals—A comprehensive survey. *IEEE Access*, 13, 18912–18936. <https://doi.org/10.1109/ACCESS.2025.3532853>
- Akshathala, S., Adnan, B., Ramesh, M., Vaidhyanathan, K., Muhammed, B., & Parthasarathy, K. (2026). Beyond task completion: An assessment framework for evaluating agentic AI systems. In *Proceedings of the 2026 International Workshop on Agentic Engineering* (pp. 9–17). ACM. <https://doi.org/10.1145/3786167.3788414>
- Alva, L. R., & Pandey, B. (2026). Agentic AI systems in the age of generative models: Architectures, cloud scalability, and real-world applications. *Artificial Intelligence Review*, 59(3), 88. <https://doi.org/10.1007/s10462-025-11458-6>
- Badiu, I., Anagnoste, S., & Savastano, M. (2025). Mapping the impact of AI agents on business performance: A bibliometric approach. *Proceedings of the International Conference on Business Excellence*, 19(1), 3492–3510. <https://doi.org/10.2478/picbe-2025-0267>
- Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., & Vilipala, S. V. (2025). The rise of agentic AI: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges. *Future Internet*, 17(9), 404. <https://doi.org/10.3390/fi17090404>
- Becattini, M., Verdecchia, R., & Vicario, E. (2025). SALLMA: A software architecture for LLM-based multi-agent systems. In *2025 IEEE/ACM International Workshop New Trends in Software Architecture (SATrends)* (pp. 5–8). IEEE. <https://doi.org/10.1109/SATrends66715.2025.00006>
- Brohi, S., Mastoi, Q., Jhanjhi, N. Z., & Pillai, T. R. (2025). A research landscape of agentic AI and large language models: Applications, challenges and future directions. *Algorithms*, 18(8), 499. <https://doi.org/10.3390/a18080499>

- Casini, S., Ducange, P., Marcelloni, F., & Pollini, L. (2025). Human-centered AI and autonomy in robotics: Insights from a bibliometric study. In *2025 International Joint Conference on Neural Networks (IJCNN)* (pp. 1–8). IEEE. <https://doi.org/10.1109/IJCNN64981.2025.11227890>
- Chowa, S. S., Alvi, R., Rahman, S. S., Rahman, M. A., Raiaan, M. A. K., Islam, M. R., Hussain, M., & Azam, S. (2026). From language to action: A review of large language models as autonomous agents and tool users. *Artificial Intelligence Review*, 59(2), 71. <https://doi.org/10.1007/s10462-025-11471-9>
- Deshmukh, S., Kuri, M., Vanarote, V., Nimbalkar, M., Chandre, P., & Yadav, S. (2025). Agentic AI: A survey of autonomous agents, architectures, and emerging applications. In *2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG)* (pp. 1–8). IEEE. <https://doi.org/10.1109/ICTBIG68706.2025.11323264>
- Ferrag, M. A., Tihanyi, N., & Debbah, M. (2026). From LLM reasoning to autonomous AI agents: A comprehensive review. *IEEE Access*, 14, 84237–84285. <https://doi.org/10.1109/ACCESS.2026.3698694>
- Gonzalez, G. R., Habel, J., & Hunter, G. K. (2026). AI agents, agentic AI, and the future of sales. *Journal of Business Research*, 202, 115799. <https://doi.org/10.1016/j.jbusres.2025.115799>
- Gridach, M., Nanavati, J., Zine El Abidine, K., Yacoubian, C., & Mack, C. (2026). Agentic AI in healthcare: Opportunities, challenges, and future directions. *ACM Computing Surveys*, 58(13), 1–30. <https://doi.org/10.1145/3809164>
- Haidemariam, T. (2026). From the logic of coordination to goal-directed reasoning: The agentic turn in artificial intelligence. *Frontiers in Artificial Intelligence*, 8, 1728738. <https://doi.org/10.3389/frai.2025.1728738>
- Hasan, M. M., Li, H., Fallahzadeh, E., Rajbahadur, G. K., Adams, B., & Hassan, A. E. (2026). An empirical study of testing practices in open source AI agent frameworks and agentic applications. *Empirical Software Engineering*, 31(5), 124. <https://doi.org/10.1007/s10664-026-10857-9>
- Hosseini, S., & Seilani, H. (2025). The role of agentic AI in shaping a smart future: A systematic review. *Array*, 26, 100399. <https://doi.org/10.1016/j.array.2025.100399>
- Karunanayake, N. (2025). Next-generation agentic AI for transforming healthcare. *Informatics and Health*, 2(2), 73–83. <https://doi.org/10.1016/j.infoh.2025.03.001>
- Khalid, O., Farooqi, A. U. H., & Bilal, M. (2026). Agentic AI: A review, applications and open research challenges. *Computing*, 108(6), 80. <https://doi.org/10.1007/s00607-026-01676-3>
- Le, H. V. (2026). Mapping the evolution of AI in tourism: Bibliometric insights and future research frontiers in the age of agentic AI. *Journal of Smart Tourism*, 6(3), 247–261. <https://doi.org/10.1177/27652157261438245>
- Li, D., & Zhu, W. (2025). Bibliometric insights into large language model-driven intelligent agents: Themes, trends and prospects. In *2025 6th International Conference on Artificial Intelligence and Computer Engineering (ICAICE)* (pp. 670–678). IEEE. <https://doi.org/10.1109/ICAICE68195.2025.11382423>
- Li, X. (2025). Task planning and decision-making methods for intelligent agents based on large language models. In *Proceedings of the 4th International Conference on Artificial Intelligence and Intelligent Information Processing* (pp. 817–822). ACM. <https://doi.org/10.1145/3778534.3778661>
- Li, X., Wang, S., Zeng, S., Wu, Y., & Yang, Y. (2024). A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges. *Vicinagearth*, 1(1), 9. <https://doi.org/10.1007/s44336-024-00009-2>
- Morales, C., Costa, L. A. L. F. D., Rigo, S. J., Kunst, R., Souza, V. C. D., Silva, É. P., Prado, G. L. E., Schardosim, T. D. C., & Roehrs, A. (2026). A systematic literature review of

- agentic AI: Definitions, architectures, and challenges. *IEEE Access*, 14, 36176–36190. <https://doi.org/10.1109/ACCESS.2026.3668138>
- Olujimi, P. A., Owolawi, P. A., Pretorius, A., & Van Wyk, E. (2025). Mapping the research landscape of agentic AI in SMEs through a bibliometric analysis of patterns and knowledge gaps. *Discover Artificial Intelligence*, 6(1), 63. <https://doi.org/10.1007/s44163-025-00764-1>
- Otoun, N., & Elkhaili, N. (2026). Methods and techniques of agentic software engineering: A systematic literature review. *IEEE Access*, 14, 7443–7465. <https://doi.org/10.1109/ACCESS.2026.3652325>
- Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (pp. 1–22). ACM. <https://doi.org/10.1145/3586183.3606763>
- Pati, A. K. (2025). Agentic AI: A comprehensive survey of technologies, applications, and societal implications. *IEEE Access*, 13, 151824–151837. <https://doi.org/10.1109/ACCESS.2025.3585609>
- Ramos, M. C., Collison, C. J., & White, A. D. (2025). A review of large language models and autonomous agents in chemistry. *Chemical Science*, 16(6), 2514–2572. <https://doi.org/10.1039/D4SC03921A>
- Raptis, E. K., Kapoutsis, A. Ch., & Kosmatopoulos, E. B. (2025). Agentic LLM-based robotic systems for real-world applications: A review on their agenticity and ethics. *Frontiers in Robotics and AI*, 12, 1605405. <https://doi.org/10.3389/frobt.2025.1605405>
- Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2026). AI agents vs. agentic AI: A conceptual taxonomy, applications and challenges. *Information Fusion*, 126, 103599. <https://doi.org/10.1016/j.inffus.2025.103599>
- Shinn, N., Cassano, F., Gopinath, A., Narasimhan, K., & Yao, S. (2023). Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 8634–8652. <https://doi.org/10.52202/075280-0377>
- Stübinger, J., & Metz, F. (2026). Understanding AI agents—A data-driven literature review. *Mathematics*, 14(9), 1478. <https://doi.org/10.3390/math14091478>
- Ugur, S. (2026). A bibliometric analysis of agentic AI and the agentic pedagogical agency framework. *International Journal of Educational Technology and Learning*, 20(1), 1–18. <https://doi.org/10.55217/101.v20i1.1048>
- Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J. (2024). A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6), 186345. <https://doi.org/10.1007/s11704-024-40231-1>
- Xu, A., Du, M., Yu, T., Puvvadi, M., Yu, T., Guo, Y., Chen, X., & Gottschlich, J. (2025a). Agentic AI for enterprise: Emerging applications and real-world challenges. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2* (pp. 6300–6301). ACM. <https://doi.org/10.1145/3711896.3737871>
- Xu, Y., Qin, M., Fu, K., Batumalay, M., Chan, C. K., Wider, W., Jiang, L., & Peng, J. (2025b). A bibliometric analysis on formation control for multi-agent system. *Social Network Analysis and Mining*, 15(1), 79. <https://doi.org/10.1007/s13278-025-01487-y>
- Zhang, B., & Zhang, C. (2026). The age of autonomous agents: A bibliometric review of agentic AI architectures, applications, and emerging challenges. *Journal of AI Analytics and Applications*, 2(2), 32–46. <https://doi.org/10.63646/KPQM1958>
- Zhu, K., Du, H., Hong, Z., Yang, X., Guo, S., Wang, Z., Wang, Z., Qian, C., Tang, X., Ji, H., & You, J. (2025). MultiAgentBench: Evaluating the collaboration and competition of LLM agents. In W. Che, J. Nabende, E. Shutova, & M. T. Pilehvar (Eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1:*

Long Papers) (pp. 8580–8622). Association for Computational Linguistics.
<https://doi.org/10.18653/v1/2025.acl-long.421>

Zrigui, F., Cheniti-Belcadhi, L., Guyeux, C., & Laiymani, D. (2026). Generative AI agents in personalized learning and assessment: A systematic bibliometric review. In *2026 IEEE International Conference on Smart Sustainable Systems for Computer and Engineering Applications (3SCEA)* (pp. 1294–1303). IEEE.
<https://doi.org/10.1109/3SCEA68071.2026.11603193>