

BRAIN. Broad Research in Artificial Intelligence and Neuroscience

e-ISSN: 2067-3957 | p-ISSN: 2068-0473

Covered in: Web of Science (ESCI); EBSCO; JERIH PLUS (hkdir.no); IndexCopernicus; Google Scholar; SHERPA/RoMEO; ArticleReach Direct; WorldCat; CrossRef; Peeref; Bridge of Knowledge (mostwiedzy.pl); abcdindex.com; Editage; Ingenta Connect Publication; OALib; scite.ai; Scholar9; Scientific and Technical Information Portal; FID Move; ADVANCED SCIENCES INDEX (European Science Evaluation Centre, neredataltics.org); ivySCI; exaly.com; Journal Selector Tool (letpub.com); Citefactor.org; fatcat!; ZDB catalogue; Catalogue SUDOC (abes.fr); OpenAlex; Wikidata; The ISSN Portal; Socolar; KVK-Volltitel (kit.edu) 2026, Volume 17, Issue 1, pages: 156-169

Submitted: January 30th, 2026 | Accepted for publication: March 3rd, 2026

The Post-ChatGPT Research Landscape of Generative Artificial Intelligence in Higher Education (2023–2026): A Bibliometric Analysis

Monica Pătruț

Vasile Alecsandri University of Bacau,
Romania

monicapatrut@yahoo.com

<https://orcid.org/0009-0007-9694-2949>

Abstract: *The launch of ChatGPT at the end of 2022 significantly accelerated research on generative artificial intelligence (GenAI) in higher education. This study conducts a bibliometric analysis of publications indexed in Web of Science during the period 2023–1 February 2026, using VOSviewer to identify dominant thematic structures, influential actors, and emerging institutional networks. The results indicate a field that remains predominantly technologically oriented, organised around three major clusters: large language models and machine learning, computational modelling and optimisation, and the educational dimension, centred on academic integrity and AI literacy. Although pedagogical interest is increasing, conceptual centrality remains concentrated around technological infrastructure. The institutional analysis highlights the role of elite universities in the United States, Asia, and Europe. The period 2023–2024 can be interpreted as a transition from technological validation and ethical concerns towards institutional and curricular integration. Overall, the study provides a structural perspective on the early post-ChatGPT phase and outlines future directions concerning sustainability and AI-assisted educational transformation.*

Keywords: *generative artificial intelligence; higher education; ChatGPT; bibliometric analysis; large language models; AI in education; research landscape.*

How to cite: Pătruț, M. (2026). The post-ChatGPT research landscape of generative artificial intelligence in higher education (2023–2026): A bibliometric analysis. *BRAIN. Broad Research in Artificial Intelligence and Neuroscience*, 17(1), 156-169. <https://doi.org/10.70594/brain/17.1/12>

1. Introduction

The rapid emergence of generative artificial intelligence (GenAI) represents one of the most significant technological disruptions in contemporary education. The public launch of ChatGPT in November 2022 marked a turning point, triggering unprecedented academic, institutional, and societal debates about the role of large language models (LLMs) in knowledge production and learning processes. Within only a few months, generative AI systems moved from experimental tools to widely used educational technologies, sparking urgent discussions about academic integrity, the redesign of assessment, pedagogical transformation, and institutional governance (Dwivedi et al., 2023; Cotton, Cotton, & Shipway, 2023).

In higher education, the integration of GenAI has proven particularly complex. Early studies focused predominantly on performance validation and the comparative evaluation of models, especially in high-stakes fields, such as medicine. Research has demonstrated that LLMs can achieve competitive results in medical examinations and clinical reasoning tasks (Thirunavukarasu et al., 2023; Gilson et al., 2023), thereby strengthening confidence in the technical robustness of these systems. At the same time, systematic reviews and rapid analyses highlighted both the promising prospects and significant risks associated with the integration of ChatGPT into academic environments (Sallam, 2023; Lo, 2023). These early contributions framed GenAI not merely as a technological innovation, but as a transformative force with profound epistemological and ethical implications.

As the scientific discourse matured, attention shifted towards issues specific to higher education. Themes such as academic integrity, student perceptions, AI literacy, and technology acceptance became central (Chan & Hu, 2023; Cotton, Cotton, & Shipway, 2023). Higher education institutions were compelled to adapt their policies, attempting to balance innovation with the maintenance of academic standards. Beyond the immediate ethical debates, structural questions emerged: How should curricula be adapted in a context of human–AI co-authorship? Which forms of assessment remain relevant in the era of automated content generation? How can universities sustainably integrate GenAI without amplifying existing digital inequalities?

Within this framework, sustainability becomes an essential dimension of analysis. The sustainability of generative artificial intelligence in higher education extends beyond ecological aspects—such as the energy and computational costs associated with training models—and encompasses institutional sustainability (long-term governance frameworks), pedagogical sustainability (the development of AI literacy competences and adaptive curriculum design), and epistemic sustainability (equitable participation in AI-assisted knowledge production). The concentration of AI infrastructure in elite universities and technologically advanced countries raises concerns regarding the epistemic asymmetries and the marginalisation of certain institutions.

In the context of the rapid growth of scientific output beginning in 2023, several bibliometric studies have attempted to map the broader research landscape of artificial intelligence in education (Kavitha & Joshith, 2024; Ng & Ho, 2025; Nimerenco & Badea, 2024). Although these analyses provide valuable insights into long-term trends and emerging clusters, most rely on extended time frames or examine artificial intelligence in education in a general manner, without clearly delineating the distinct post-ChatGPT phase. Furthermore, the existing literature gives limited attention to the integrated analysis of thematic structures, institutional ecosystems, and influential actors during the critical period of 2023–2026. The years immediately following the launch of ChatGPT represent a compressed phase of scientific acceleration, characterised by the sharp growth in citations, the rapid consolidation of research themes, and the emergence of global collaboration networks that developed at an exceptionally rapid pace. Understanding this early configuration is essential for anticipating the future trajectory of GenAI in higher education and identifying structural imbalances or research gaps that may shape the field's development.

Therefore, this study aims to conduct a comprehensive bibliometric analysis of research on generative artificial intelligence in higher education from 2023 to 1 February 2026, using data from

the Web of Science and visualisation and mapping techniques implemented through VOSviewer. The specific objectives of the study are:

- To identify the dominant thematic clusters and analyse their evolution within the post-ChatGPT landscape;
- To analyse the authors, institutions and journals exerting major influence in the field;
- To conceptualise the structural configuration of research on GenAI in higher education and to evaluate its implications for pedagogical integration and sustainability.

By focusing on the immediate post-disruption period, the study makes three main contributions. First, it delineates generative artificial intelligence as a distinct wave within the broader field of AI & Education research. Second, it integrates thematic, institutional and citation analyses in order to provide a multidimensional perspective on the field. Third, it proposes an interpretative model of the transition from the phase of technological validation to pedagogical integration for the 2023–2024 period.

Thus, the study provides both a diagnostic mapping of the current state of the field and a strategic reflection on the future directions of generative artificial intelligence in higher education.

2. Methods

The search strategy was designed to specifically select recent literature concerning the use of generative artificial intelligence in higher education. The query was formulated in the **Topic (TS)** field of the Web of Science Core Collection database, combining terms referring to generative technologies (“generative AI”, “generative artificial intelligence”, “large language model*”, “ChatGPT”, “GPT-3”, “GPT-4”) with terms specifying the targeted educational context (“higher education”, “universit*”, “tertiary education”, “college student*”), using the logical operator **AND**. Only documents of the types “Article”, “Review Article”, and “Early Access” were included, ensuring a focus on primary scientific output and synthesis publications published in English during the period when GenAI technologies became widely available and discussed in the academic environment.

Given that the public launch of ChatGPT occurred at the very end of 2022 and that scientific publishing involves inherent delays, we selected 2023 as the starting year for article inclusion, as it marks the year in which the first wave of peer-reviewed publications appeared. Moreover, following common practice in bibliometric reviews, we restricted our corpus to peer-reviewed journal articles and review papers indexed in Web of Science. Conference papers were excluded due to their heterogeneous review processes, their generally lower and more variable citation impact, and the risk of duplicating content later published as full journal articles. This decision was intended to ensure a more consistent level of quality and comparability across documents.

To clearly illustrate the study selection process, the steps described above were synthesised in a **PRISMA-type flow diagram** (see Fig. 1), in accordance with the **PRISMA 2020** recommendations for transparent reporting of the record inclusion process and their adaptation to bibliometric analyses (Page et al., 2021).

In the first stage, a total of **36,483 records** were identified in the Web of Science database using the predefined search strategy, which combined a **Topic (TS)** query with a temporal limitation to the period **2023–1 February 2026**, the document types “**Article**”, “**Review**”, and “**Early Access**”, and a restriction to publications in **English**.

In the second stage, the **36,483 records** underwent a systematic screening process based on document type and relevance criteria aligned with the study’s objective. During this stage, **15,828 records** were excluded, mainly books, book chapters, conference proceedings papers, and retracted publications that did not correspond to the proposed bibliometric design.




Figure 1. PRISMA 2020 Flow Diagram

After applying the inclusion and exclusion criteria, **20,655 records** met the eligibility criteria for inclusion in the third stage and were retained for further analysis. The eligible corpus consisted exclusively of documents classified as “**Article**”, “**Review Article**”, and “**Early Access**”, to ensure the comparability and quality of the analysed data.

In the final stage, the **20,655 eligible studies** were included in the bibliometric analysis itself. The corpus comprised **16,961 original research articles (Articles)**, **1,921 publications classified as “Early Access”**, and **1,773 review articles (Review Articles)**, providing a comprehensive overview of recent scientific output in the investigated field.

3. Results

3.1. Publication Growth Over Time

The **temporal distribution of publications** shows a trajectory of rapid growth following the launch of ChatGPT at the end of 2022. The number of publications increased substantially from **1,452 in 2023** to **5,813 in 2024**, then rose further to **12,596 in 2025**.

Moreover, in the period **1 January – 1 February 2026**, **794 publications** had already been indexed, indicating sustained momentum in research on generative artificial intelligence in higher education (see Figure 2).




Figure 2. Distribution of Publications by Year

3.2. Geographic Distribution of Research Output

The analysis of geographic distribution conducted using **VOSviewer** shows a strong concentration of research on generative artificial intelligence in higher education in a limited number of countries with well-developed academic and technological infrastructures. Of the **162 countries** identified in the corpus, **104 exceed the threshold of at least 5 documents and 5 citations**, while the **top 25 countries** have the highest values for both publication output and citation counts. The **United States (6,031 documents, 79,942 citations)** and **China (4,556 documents, 45,405 citations)** clearly dominate the field, followed by the **United Kingdom, Germany, Italy, Australia, and Canada**, confirming trends identified in other studies on AI in education.

At the same time, countries in **Asia and the Gulf region** (Saudi Arabia, the United Arab Emirates, Singapore, South Korea, India, and Malaysia) show increasing publication volumes and collaboration intensity, indicating that GenAI in higher education is becoming a strategic research topic amid accelerated educational reform.




Figure 3. Country collaboration network in post-ChatGPT GenAI research (2023–2026)

The **international collaboration network** (see Figure 3) also reflects a densely interconnected global structure in which the **United States occupies a structurally central position**, both through its high volume of publications and the intensity of its collaborative links. Strong connectivity with **Europe, Asia, and Oceania** indicates a consolidated pattern of transcontinental cooperation. At the same time, there is a strong presence of Asian countries (China, Singapore, South Korea, India, and Japan), confirming the accelerated dynamics of research ecosystems in this region. Singapore and South Korea appear as actors with a high level of integration into global research networks, while Germany and the Netherlands reinforce the European pole of scientific collaboration.

This distribution also indicates a **rapid and almost simultaneous mobilisation of the global academic community** following the launch of ChatGPT. Overall, the figure confirms the deeply **internationalised nature of the field** and reveals a **multipolar structure**, in which **intercontinental collaboration plays a crucial role** in shaping the early post-ChatGPT research landscape.

3.3. Most Influential Publications

Table 1. Most Cited Publications on Generative AI in Higher Education

Rank	Publications	Year	Citations
1	Dwivedi, Y. K., Kshetri, N., Hughes, L., et al. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. <i>International Journal of Information Management</i> , 71, 102642.	2023	2008
2	Thirunavukarasu, A. J., Ting, D. S. J., Elangovan, K., Gutierrez, L., Tan, T. F., & Ting, D. S. W. (2023). Large language models in medicine. <i>Nature Medicine</i> , 29, 1930–1940.	2023	1921
3	Sallam, M. (2023). ChatGPT utility in healthcare education, research, and practice: Systematic review on the promising perspectives and valid concerns. <i>Healthcare</i> , 11(6), 887.	2023	1598
4	Ayers, J. W., Poliak, A., Dredze, M., et al. (2023). Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. <i>JAMA Internal Medicine</i> .	2023	1392
5	Gilson, A., Safranek, C. W., Huang, T., et al. (2023). How does ChatGPT perform on the United States Medical Licensing Examination (USMLE)? The implications of large language models for medical education and knowledge assessment. <i>JMIR Medical Education</i> .	2023	1271
6	Chang, Y., Wang, X., Wang, J., et al. (2024). A survey on evaluation of large language models. <i>ACM Transactions on Intelligent Systems and Technology</i> , 15(3), Article 39.	2024	1268
7	Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2023). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. <i>Innovations in Education and Teaching International</i> , 61(2), 228–239.	2023	1061
8	Wu, T., He, S., Liu, J., et al. (2023). A brief overview of ChatGPT: The history, status quo and potential future development. <i>IEEE/CAA Journal of Automatica Sinica</i> , 10(5).	2023	873
9	Lo, C. K. (2023). What is the impact of ChatGPT on education? A rapid review of the literature. <i>Education Sciences</i> , 13(4).	2023	872
10	Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. <i>International Journal of Educational Technology in Higher Education</i> , 20, Article 43.	2023	837

The analysis of the **ten most cited articles** on generative artificial intelligence highlights a significant concentration of scientific impact in **2023**. Most of the publications in the ranking were published in that year, reflecting a rapid **“citation burst”** and an accelerated mobilisation of the academic community. The article coordinated by Dwivedi et al. (2023) ranks first with **over 2,000**

citations and serves as a major **interdisciplinary reference work** that synthesises the opportunities, challenges, and implications of **ChatGPT-like technologies** for research, practice, and educational policy.

A notable aspect is the **predominance of the medical field** among the most highly cited articles (e.g., *Nature Medicine*, *JAMA Internal Medicine*, *JMIR Medical Education*, *Healthcare*). These studies examine the performance of **large language models (LLMs)** in medicine, the use of **ChatGPT in medical education**, and its potential role in supporting **diagnosis and patient communication**.

Subsequently, the **educational literature** begins to gain visibility through articles such as those by Cotton, Cotton, and Shipway (2023) and Chan and Hu (2023), indicating a gradual shift towards **pedagogical integration** and reflection on **academic integrity, assessment, and student perceptions**.

Overall, **Table 1** reflects an **early consolidation phase** of the scientific literature, initially dominated by **technological validation and medical applications**, followed by the **emergence of a more specific educational research agenda**.

3.4. Leading Journals

The analysis of the **ten most productive journals** (see Table 2) reveals a **dual structure** within the editorial landscape, characterised by significant differences between publication volume and the relative impact of those publications. Although IEEE Access ranks first in the number of published articles ($n = 412$), its average impact (6.96 citations per article) is moderate compared with other journals in the ranking. Similarly, Scientific Reports and Applied Sciences-Basel present high publication volumes, though with a relatively balanced impact.

In contrast, journals specialising in the **educational and interdisciplinary domains** demonstrate a significantly higher average citation impact. Computers and Education: Artificial Intelligence records the highest average number of citations per article (25.58), followed by NPJ Digital Medicine (22.65) and Education and Information Technologies (19.24). This distribution suggests that, although journals with a technological orientation publish a larger volume of studies, journals focused on digital education and interdisciplinary applications generate a comparatively higher scientific impact.

Table 2. Top 10 Journals with the Highest Number of Publications on Generative AI in Higher Education

Rank	Journal	Publications	Citations	Citations per Article	Total Link Strength (TLS)
1	IEEE Access	412	2868	6.96	132146
2	Scientific Reports	287	2306	8.03	75864
3	Applied Sciences-Basel	239	1569	6.56	84897
4	Education and Information Technologies	181	3482	19.24	50749
5	Electronics	143	634	4.43	52887
6	Computers and Education: Artificial Intelligence	130	3326	25.58	48069
7	Information	122	867	7.11	46327
8	Expert Systems with Applications	120	693	5.78	57889
9	NPJ Digital Medicine	113	2560	22.65	38758
10	Information Processing & Management	80	355	4.44	47226

The **Total Link Strength (TLS)** also indicates a high level of interconnectivity for journals such as IEEE Access and Applied Sciences-Basel, reflecting their central position within the co-citation network.

Overall, the results confirm the predominantly technological character of the field and highlight the consolidation of several editorial hubs—including Computers and Education: Artificial Intelligence, Education and Information Technologies, and Information—with significant educational impact.

3.5. Most Productive Authors

At the **micro level**, the analysis of author and institutional productivity indicates a further concentration of expertise within several academic hubs. Among the more than 72,000 authors included in the corpus, 1,247 exceed the threshold of at least five articles and five citations. The ranking is dominated by researchers affiliated with universities in East Asia and North America, such as Nanyang Technological University, University of Hong Kong, Guangdong University of Technology, Yale University, and the University of Georgia. Out of more than 13,700 organisations, 1,827 surpass similar thresholds. At the same time, elite institutions such as Stanford University, the University of Hong Kong, the Chinese Academy of Sciences, the University of Toronto, Cambridge, and Oxford appear among the top positions, confirming the results of other bibliometric studies that highlight the central role of universities from the United States, China, the United Kingdom, and Canada in research on AI in education (Table 3).

Table 3. Top 10 Most Productive Authors in Generative AI Research (Higher Education)

Rank	Author	Institution	Country	Publications	Citations	Citations per Article
1	Dusit (Tao) Niyato	Nanyang Technological University	Singapore	71	1891	26.63
2	Hongyang Du	University of Hong Kong	Hong Kong	55	1517	27.58
3	Jiawen Kang	Guangdong University of Technology	China	45	1465	32.56
4	Zehui Xiong	Queen's University Belfast	UK	36	1018	28.28
5	Xuemin (Sherman) Shen	University of Waterloo	Canada	24	842	35.08
6	Tianming Liu	University of Georgia	USA	24	692	28.83
7	Xiang Lisa Li	Stanford University	USA	24	664	27.67
8	Dong In Kim	Sungkyunkwan University (SKKU)	South Korea	23	827	35.96
9	Zhengliang Liu	University of Georgia	USA	21	646	30.76
10	David Chartash	Yale University	USA	11	1332	121.09

At the same micro level, the analysis of the **most productive authors** reveals an interesting combination of publication volume and relative impact. Dusit (Tao) Niyato leads the ranking with 71 publications and 1,891 citations, confirming a consistent and influential presence in the fields of technological infrastructure and modelling.

Hongyang Du and Jiawen Kang complete the Asian core of high productivity, reflecting the consolidation of research ecosystems in Asia. A notable aspect concerns authors with a high relative impact in relation to their number of publications. David Chartash, for example, records 1,332 citations for 11 articles (121.09 citations per article), indicating an extremely influential contribution to the early literature on LLMs and medical applications. This distribution suggests a polarisation between authors oriented towards publication volume and those whose impact is concentrated in seminal contributions.

3.6. Institutional Contribution and Collaboration

Out of the 13,700 institutions included in the corpus, 1,827 exceeded the threshold of at least five articles per institution and five citations per institution. The institutional ranking highlights a strong concentration of scientific output within universities and research centres in the United States, Asia, and Europe. The University of Hong Kong and the Chinese Academy of Sciences lead in publication volume, while Stanford University stands out for its high total impact (8,011 citations) and a consistently high average number of citations per article (32.82) (Table 4).

Table 4. Top 10 Institutions by Publication Output in Generative AI Research (Higher Education)

Rank	Institution	Country	Publications	Citations	Citations per Article	Total Link Strength (TLS)
1	University of Hong Kong	Hong Kong	320	5587	17.46	151
2	Chinese Academy of Sciences	China	313	4390	14.03	301
3	Stanford University	USA	244	8011	32.83	157
4	University of Toronto	Canada	237	3908	16.49	28
5	Harvard Medical School	USA	218	3363	15.43	54
6	Nanyang Technological University	Singapore	214	3158	14.76	78
7	University of Oxford	UK	214	2908	13.59	57
8	University of Illinois	USA	213	3453	16.21	61
9	The Chinese University of Hong Kong	Hong Kong	176	3428	19.48	113
10	University of Cambridge	UK	173	4319	24.97	84

The **Citations per Article** indicator reveals structural differences between institutions oriented towards publication volume and those characterised by a higher relative impact. Universities such as the University of Cambridge and Stanford University produce scientific articles with particularly high global visibility.

At the same time, Total Link Strength (TLS) values indicate a high level of interconnectivity for institutions such as the Chinese Academy of Sciences, the University of Hong Kong, and Stanford University, confirming their central role within the collaboration network. Overall, the results support the idea of a field dominated by consolidated technological-institutional ecosystems, in which scientific impact is generated both through high publication volume and through seminal contributions produced during the early post-ChatGPT phase.

3.7. Thematic Structure (Keyword Clusters)

The **co-occurrence map of terms** identified four major clusters, outlining both the educational core of research on GenAI in higher education and several related technical and disciplinary areas. Figure 4 illustrates the thematic structure of the field through a keyword co-occurrence analysis (minimum threshold of 150 terms), highlighting a densely interconnected network organised around several dominant conceptual hubs.

The terms “artificial intelligence” and “large language models” occupy central positions both in terms of node size (high frequency) and the number of connections, confirming the role of technological infrastructure as the central pivot of the post-ChatGPT literature.




Figure 4. Co-occurrence map of keywords on generative artificial intelligence in higher education.

The **co-occurrence analysis of terms** generated **four major thematic clusters** that structure the research landscape on generative artificial intelligence in higher education.

Cluster 1 (57 terms) concentrates the **educational core** of the field and is dominated by keywords such as *higher education, students, education, engagement, educational technology, feedback, assessment, academic integrity, ethics, AI literacy, technology acceptance, technology adoption, intention, and trust*, alongside general terms such as *ChatGPT, GenAI, and generative AI*. This cluster reflects the intersection of **three central subfields**: **teaching and learning** (through references to engagement, feedback, and student performance), **assessment and academic integrity** (*assessment, academic integrity, ethics, trust*), and **AI literacy and curriculum development** (*AI literacy, technology acceptance/adoption, intention*). Overall, this cluster suggests that most studies focus on **pedagogical impact, perceptions, and behavioural responses** to GenAI within the university environment.

Cluster 2 (49 terms) primarily captures the **technical and methodological dimension** of the field, including terms such as *large language model(s)*, *chatbots*, *foundation models*, *transformers*, *diffusion models*, *data mining*, *predictive models*, *feature extraction*, and *sentiment analysis*, as well as *safety*, *security*, *privacy*, *costs*, *resource management*, *systematic literature review*, and *surveys*. This cluster can be interpreted as referring to the **technological infrastructure required for GenAI use in higher education**. It is closely related to the subfield of **GenAI as a research tool** (through the emphasis on predictive models, data mining, and analytical methods) and, to some extent, to issues of **sustainability and equity/infrastructure** (through terms associated with costs, resource management, security, and privacy). This indicates that a portion of the literature focuses on **technological architecture and model performance evaluation**, rather than on explicit pedagogical scenarios.

Cluster 3 (39 terms) outlines a **well-defined subfield centred on medical education and healthcare applications** of GenAI. It includes terms such as *health, healthcare, digital health, cancer, depression, mental health, diagnosis, risk, GPT, GPT-4, LLM(s), medical education, natural*

language processing, neural networks, and prompt engineering. This cluster shows that a significant portion of research on GenAI in higher education is concentrated in the **medical domain**, where generative models are used both for **training medical students** and for **clinical practice and decision support systems**. Consequently, this cluster intersects several subfields—**teaching and learning, assessment, and ethics**—but within a **discipline-specific context**.

Cluster 4 includes five terms (*information, internet, patient education, quality, and readability*) and represents a **small cluster focused on the quality and readability of information**, particularly in the context of **patient education and online resources**. Although peripheral to the broader educational core, this cluster indicates emerging concerns regarding the **quality of GenAI-generated content** and the **clarity of information**, with implications for **sustainability and equity** (equitable access to understandable information) as well as **policy and governance** (institutional responsibility for monitoring the quality of digital information).

Overall, the results show that the subfields of teaching and learning, assessment and academic integrity, and AI literacy and curriculum development are well represented within the dominant clusters. In contrast, themes explicitly related to policy, governance, sustainability and equity in higher education appear marginal or only indirectly represented (primarily through technical terms associated with costs, security, and information quality). This distribution supports the hypothesis that the literature is concentrated on the pedagogical and psychological dimensions of GenAI in higher education, while studies addressing policy frameworks, institutional governance, and the impact on educational equity and sustainability remain underrepresented.

4. Discussion

Compared with recent bibliometric studies devoted to AI and ChatGPT in education, the present analysis consolidates a **methodological direction that is still in the process of stabilisation**. Most previous research has relied on the Scopus database and has addressed either AI in higher education in general terms (Kavitha & Joshith, 2024; López-Chila et al., 2024), ChatGPT in specific educational contexts (Ivanova, Grosseck, & Holotescu, 2024; Polat et al., 2024; Koo, 2025), or has combined bibliometric approaches with thematic modelling or critical discourse analyses (Lachheb et al., 2025; Sahar & Munawaroh, 2025). Only a limited number of studies have explicitly examined generative artificial intelligence in higher education based on Web of Science data within a clearly delimited post-ChatGPT time frame.

Consistent with previously reported findings (Ng & Ho, 2025), the present study confirms the dominant themes related to academic integrity, ethics, and technological adoption. However, the co-occurrence network analysis conducted here also reveals the persistence of a strong technological centrality, suggesting that the pedagogical discourse continues to evolve within a framework still dominated by LLM infrastructure. The semantic centrality of terms such as *artificial intelligence, large language models, and machine learning* confirms that technological infrastructure constituted the initial foundation of the scientific discourse. At the same time, the emergence of themes related to academic integrity, AI literacy, and institutional adoption indicates a gradual shift towards pedagogical integration. This transition suggests that higher education did not initially generate the theoretical framework for GenAI, but rather responded to external technological developments, progressively adapting them to its own curricular and institutional structures.

The results indicate that the period 2023–1 January 2026 may be interpreted as a phase of accelerated stabilisation of research on generative artificial intelligence in higher education. The rapid growth in the number of publications immediately following the launch of ChatGPT points to a swift mobilisation of the academic community.

The geographical and institutional distribution reinforces this interpretation, highlighting the dominant role of mature technological ecosystems in the United States and Asia, while European contributions more frequently focus on ethical and pedagogical dimensions. Overall, the field

appears to be undergoing a transition from validation and experimentation towards institutional consolidation and strategic reflection.

The present analysis contributes to the emerging literature on GenAI in higher education by identifying a two-stage structural dynamic: an initial phase of technological anchoring, followed by a stage of pedagogical and institutional expansion. This trajectory confirms recent observations regarding the predominantly technological orientation of the early literature on AI in education (Kavitha & Joshith, 2024; Ivanova, Grosseck, & Holotescu, 2024; Ng & Ho, 2025), while also suggesting a progressive reconfiguration of thematic priorities after the launch of ChatGPT.

The semantic centrality of concepts associated with large language models and machine learning indicates that the academic legitimacy of GenAI was initially negotiated through technological performance and validation, a phenomenon also observed in other bibliometric analyses of ChatGPT (Polat et al., 2024; Koo, 2025). However, the simultaneous emergence of themes related to academic integrity and AI literacy signals a phase of pedagogical internalisation of the technology, in which higher education institutions attempt to develop their own normative and interpretative frameworks (Lachheb et al., 2025).

From an epistemological perspective, the results suggest that the field remains in a pre-paradigmatic phase, characterised by the coexistence of technological enthusiasm and critical reflection. This observation refines the conclusions of recent studies mapping GenAI in education by providing a clear temporal delimitation of the post-disruption phase generated by the launch of ChatGPT.

From a practical perspective, the findings highlight the need for a strategic and institutionalised approach to integrating GenAI into higher education. First, recent research has emphasised concerns regarding academic integrity and the responsible use of ChatGPT (Cotton, Cotton, & Shipway, 2023), and the present analysis confirms the centrality of these themes within the emerging conceptual network. Second, studies on student perceptions indicate both perceived benefits and uncertainties regarding the use of GenAI in learning processes (Chan & Hu, 2023), suggesting the need to develop AI literacy competences at the curricular level. Applications in medical and evaluative contexts have rapidly demonstrated both the potential and the limitations of LLMs in high-stakes educational settings (Gilson et al., 2023; Thirunavukarasu et al., 2023), highlighting the need to develop institutional mechanisms for regulation and continuous evaluation.

Furthermore, the geopolitical distribution of scientific output suggests the risk of reinforcing global asymmetries in access to infrastructure and expertise, which calls for reflection on the equity of participation in GenAI-assisted educational transformation and on its sustainable integration as both an institutional and ethical process.

5. Limitations and Future Research Directions

Although the bibliometric analysis provides a coherent structural overview of the evolution of research on generative artificial intelligence in higher education, the study has several limitations that should be acknowledged.

First, the analysed corpus was extracted exclusively from the **Web of Science** database and limited to publications written in English. This methodological choice ensures a high level of data quality and comparability, but it may exclude relevant contributions indexed in other international databases or published in other languages.

Second, the exclusive inclusion of documents classified as **Article**, **Review**, and **Early Access** resulted in the exclusion of conference papers and book chapters, which may be highly relevant in an emerging technologically-oriented field.

Third, the quantitative nature of bibliometric analysis restricts interpretation to relational structures (such as co-citation, co-occurrence, and collaboration networks), without allowing for an in-depth evaluation of the conceptual content or the quality of the theoretical arguments within the analysed works. Therefore, the results capture the field's structural configuration but do not provide a detailed analysis of epistemological positions or internal theoretical divergences.

Regarding **future research directions**, comparative analyses across databases (such as **Web of Science** and **Scopus**) would be valuable to assess the robustness of the identified thematic structure. The integration of **mixed-methods approaches**, combining bibliometrics with qualitative content analysis or topic modelling techniques, could also provide a more nuanced understanding of the field's conceptual evolution.

Furthermore, **longitudinal studies** extending the analysed time frame would allow researchers to observe the paradigmatic maturation of research on GenAI in higher education. Finally, investigating **regional differences in the pedagogical and institutional integration of GenAI** could contribute to a better understanding of the global implications of AI-assisted educational transformation.

6. Conclusions

The present study provides a comprehensive bibliometric analysis of research on generative artificial intelligence in higher education **from 2023 to 1 February 2026**, capturing the early structural configuration of the **post-ChatGPT research landscape**. The findings reveal an **exponential increase in scientific output** immediately following the launch of ChatGPT, confirming the disruptive character of this technological moment for the global educational ecosystem.

The thematic analysis highlights a conceptual structure organised around a **dominant technological core** (large language models, machine learning, and natural language processing), while the **pedagogical dimension**, including academic integrity and AI literacy, is progressively consolidating. This distribution suggests that higher education is currently undergoing a **transition from technological validation and experimentation towards institutional integration and strategic reflection**.

At the **geopolitical and institutional levels**, the field is shaped by well-established research ecosystems, particularly in the **United States and Asia**, indicating the presence of influential scientific hubs that shape the global research agenda. At the same time, the emergence of themes related to **ethics, governance, and sustainability** signals the gradual maturation of the debate.

By explicitly delineating the early **post-disruption phase**, the study contributes to understanding how an emerging educational field becomes structured around a transformative technology. Overall, the research provides a **reference framework for future investigations** concerning theoretical consolidation, institutional adaptation, and the sustainable integration of generative artificial intelligence in higher education.

References

- Ayers, J. W., Poliak, A., Dredze, M., Leas, E. C., Zhu, Z., Kelley, J. B., Faix, D. J., Goodman, A. M., Longhurst, C. A., Hogarth, M., & Smith, D. M. (2023). Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Internal Medicine*, 183(6), 589–596. <https://doi.org/10.1001/jamainternmed.2023.1838>
- Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20, 43. <https://doi.org/10.1186/s41239-023-00411-8>
- Chang, Y., Wang, X., Wang, J., Wu, Y., Yang, L., Zhu, K., Chen, H., Yi, X., Wang, C., Wang, Y., et al. (2024). A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 15(3), Article 39, 1–45. <https://doi.org/10.1145/3641289>
- Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2023). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 61(2), 228–239. <https://doi.org/10.1080/14703297.2023.2190148>
- Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S.,

- Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., ... Wright, R. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
- Gilson, A., Safranek, C. W., Huang, T., Socrates, V., Chi, L., Taylor, R. A., & Chartash, D. (2023). How does ChatGPT perform on the United States Medical Licensing Examination (USMLE)? The implications of large language models for medical education and knowledge assessment. *JMIR Medical Education*, 9, e45312. <https://doi.org/10.2196/45312>
- Ivanova, M., Grosseck, G., & Holotescu, C. (2024). Unveiling insights: A bibliometric analysis of artificial intelligence in teaching. *Informatics*, 11(1), 10. <https://doi.org/10.3390/informatics11010010>
- Kavitha, K., & Joshith, V. P. (2024). The transformative trajectory of artificial intelligence in education: Two decades of bibliometric retrospect. *Journal of Educational Technology Systems*, 52(3). <https://doi.org/10.1177/00472395241231815>
- Koo, M. (2025). ChatGPT research: A bibliometric analysis based on the Web of Science from 2023 to June 2024. *Knowledge*, 5(1), 4. <https://doi.org/10.3390/knowledge5010004>
- Lachheb, A., Leung, J., Abramenska-Lachheb, V., & Sankaranarayanan, R. (2025). AI in higher education: A bibliometric analysis, synthesis, and a critique of research. *The Internet and Higher Education*, 67, 101021. <https://doi.org/10.1016/j.iheduc.2025.101021>
- Lo, C. K. (2023). What is the impact of ChatGPT on education? A rapid review of the literature. *Education Sciences*, 13(4), 410. <https://doi.org/10.3390/educsci13040410>
- López-Chila, R., Llerena-Izquierdo, J., Sumba-Nacipucha, N., & Cueva-Estrada, J. (2024). Artificial intelligence in higher education: An analysis of existing bibliometrics. *Education Sciences*, 14(1), 47. <https://doi.org/10.3390/educsci14010047>
- Ng, S. L., & Ho, C. C. (2025). Generative AI in education: Mapping the research landscape through bibliometric analysis. *Information*, 16(8), 657. <https://doi.org/10.3390/info16080657>
- Nimerenco, I., & Badea, A. M. (2024). Bibliometric analysis of artificial intelligence's impact on education: A comprehensive overview. *Annals of "Dunarea de Jos" University of Galati*, 3(1), 418–432. <https://doi.org/10.35219/eai15840409470>
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
- Polat, H., Topuz, A. C., Yildiz, M., Taşlıbeyaz, E., & Kurşun, E. (2024). A bibliometric analysis of research on ChatGPT in education. *International Journal of Technology in Education*, 7(1), 59–85. <https://doi.org/10.46328/ijte.606>
- Sahar, R., & Munawaroh, M. (2025). Artificial intelligence in higher education with bibliometric and content analysis for future research agenda. *Discover Sustainability*, 6, 401. <https://doi.org/10.1007/s43621-025-01086-z>
- Sallam, M. (2023). ChatGPT utility in healthcare education, research, and practice: Systematic review on the promising perspectives and valid concerns. *Healthcare*, 11(6), 887. <https://doi.org/10.3390/healthcare11060887>
- Thirunavukarasu, A. J., Ting, D. S. J., Elangovan, K., Gutierrez, L., Tan, T. F., & Ting, D. S. W. (2023). Large language models in medicine. *Nature Medicine*, 29, 1930–1940. <https://doi.org/10.1038/s41591-023-02448-8>
- Wu, T., He, S., Liu, J., Sun, K., Liu, L., Han, L., & Tang, Y. (2023). A brief overview of ChatGPT: The history, status quo and potential future development. *IEEE/CAA Journal of Automatica Sinica*, 10(5), 1122–1136. <https://doi.org/10.1109/JAS.2023.123618>