Article

Artificial Intelligence in Education: A Bibliometric and Knowledge Diffusion Analysis Based on OpenAlex

WANG Jing [1], CHEN Hao [2,*]

[1] Vietnam National University, Hanoi, Vietnam; wangjing@vnu.edu.vn
[2] De La Salle University, Manila, Philippines; chenhao@dlsu.edu.ph
* Correspondence: chenhao@dlsu.edu.ph

Abstract

The Journal of Sustainability Insights, vol. 1, no. 1, November 1, 2026. DOI: 10.00000/tjosi. To map the research landscape, growth dynamics and knowledge-diffusion structure of artificial intelligence (AI) in education, this study draws on the OpenAlex database and adopts a multi-query strategy to collect publications from 2000 to 2025. Logistic and Gompertz growth curves are fitted to cumulative publication counts, and a negative binomial regression is estimated on highly cited papers to identify the determinants of citation intensity. The results show that AI-in-education research remained modest for more than two decades before entering explosive growth after the emergence of generative AI in 2022, with annual output rising from 657 papers in 2022 to 5,912 in 2025. The Logistic model fits marginally better ($R^2 = 0.9989$) than the Gompertz model, but its asymptotic ceiling is pushed upward by the recent surge, suggesting the field is far from saturation. Output is highly concentrated in China, the United States, India and other countries, with the top five accounting for 41.8%. Ethics, privacy and algorithmic bias have emerged as cross-cutting hotspots. Papers on the core topic of AI in education receive significantly higher citation intensity (IRR=1.476, p=0.008). The study traces a diffusion pathway from technology provision towards human-centred values and ethical governance.

Keywords: AI in education; bibliometrics; knowledge diffusion; OpenAlex; growth curves; negative binomial regression; generative AI

1. Introduction

The convergence of artificial intelligence and education is reshaping the fundamentals of teaching and learning. From early intelligent tutoring systems (ITS) and adaptive learning platforms to the wave of generative AI triggered by ChatGPT in late 2022, each technological leap has raised new

Received: 24 September 2026
Revised: 7 October 2026 Accepted: 9 October 2026 Published: 1 November 2026

Copyright: © 2026 by the authors. Licensee Lightup Innovation, United Kingdom, HR5 3DJ. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license

questions within the educational research community. Artificial intelligence in education (AIEd) has evolved from a marginal technical experiment into an independent interdisciplinary field spanning education, computer science and cognitive science. Understanding how this field grows, who drives it and how its knowledge diffuses is not only of scholarly-historical interest but also bears directly on how educational institutions allocate resources and formulate technology-adoption and governance policies.

Technological optimism and technological scepticism have long coexisted in the academic discourse on AI in education. Proponents emphasise the potential of AI for personalisation at scale, immediate feedback and reducing educational costs, regarding it as key to addressing unequal access to educational resources and realising instruction tailored to individual learners[1,2]; more cautious voices warn about algorithmic bias, data privacy, academic integrity and the erosion of human agency[3,4]. The tension between these two narratives drives the field's development and shapes the evolution of its research agenda. When ChatGPT entered public view in late 2022 through conversational interaction, this tension was amplified: educators actively explored its use in lesson preparation, writing and assessment[5,6], universities rapidly introduced policies to address cheating and academic-integrity challenges, and scholars engaged in intensive debate around 'AI literacy', 'prompt engineering' and 'critical thinking'[7]. This coexistence of demand and concern provides an ideal natural experiment for observing knowledge diffusion.

Yet confronted with an exponentially expanding literature, close reading within a few individual reviews can no longer capture the whole picture. Bibliometrics, by quantifying the production, dissemination and citation of documents, offers a replicable and verifiable methodological toolkit for identifying research fronts and mapping intellectual structures. Existing bibliometric studies of AI in education have mostly relied on a single data source—Web of Science or Scopus—and their time windows largely end before the generative-AI explosion, leaving the structural rupture since 2023 insufficiently characterised. Moreover, 'knowledge diffusion'—the process by which concepts, topics and methods migrate and become institutionalised across countries, institutions and disciplines—has not been systematically represented in these studies.

OpenAlex was chosen as the data source also for reasons of methodological reproducibility. The credibility of bibliometric research depends heavily on the transparency of data acquisition and processing. Commercial databases are often subject to licensing restrictions that prevent researchers from publishing complete search strings and raw records, making comparable studies difficult to verify. OpenAlex provides an open API and stable entity identifiers, allowing every step from retrieval through cleaning to aggregation to be re-run by others. We therefore make our search strategy and processing scripts publicly available, seeking to ground the conclusions in a traceable chain of evidence. This orientation also aligns with recent advocacy for research transparency in the open-science movement.

This study aims to fill the gap. OpenAlex was selected because it is open, free, broad in coverage, rapidly updated, and offers fine-grained fields for topics, keywords, open-access status and citation relations, facilitating reproducible macro-level bibliometrics and micro-level regression analysis. Three research questions guide the inquiry. First, what temporal patterns characterise the growth of AI-in-education output over more than two decades, and can classic growth curves describe and forecast its future trajectory? Second, how is output distributed geographically and institutionally, and does knowledge production exhibit centralisation or dispersion? Third, what features charac-

terise the intellectual base and diffusion pathways as revealed by topic structures and highly cited papers, and which factors influence the knowledge impact of individual papers, measured by citation count?

The remainder of the article is organised as follows. Section 2 reviews the development of AI-in-education research and bibliometric methodology before presenting the analytical framework and research questions. Section 3 details the data sources, search strategy and analytical methods. Section 4 reports empirical results on temporal evolution, growth fitting, spatial distribution, topic structure and impact regression. Section 5 discusses the diffusion effects of generative AI, the limits of growth forecasts, inequalities in knowledge production and the value shift, while openly acknowledging limitations. Section 6 concludes and draws implications for researchers, educational administrators and policymakers.

The marginal contributions of this study are threefold. First, the observation window is extended to 2025, capturing the discontinuous growth of research output under the shock of generative AI and quantifying its stage characteristics with growth-curve models. Second, macro-level distributions and micro-level regressions are combined, bringing questions of 'who produces knowledge' and 'what kind of knowledge attracts greater impact' into a unified analytical framework. Third, all data and analytical procedures rely on an open database and open-source tools and are fully reproducible, providing a methodological reference for similar studies.

2. Literature review and theoretical framework

2.1. The research lineage of AI in education

The lineage of AIEd research can be traced to the prototypes of intelligent tutoring systems in the 1960s and 1970s, although a sizeable academic community did not form until after the turn of the century. In their systematic review of papers published since the founding of the International Journal of Artificial Intelligence in Education, Roll and Wylie characterised the field's evolution as two parallel paths—'evolution' and 'revolution': the former focuses on classroom practice, collaboration with teachers and the diversification of technologies and disciplines, while the latter advocates embedding technology in learners' everyday lives and supporting their cultural practices and learning communities[8]. This dichotomy set the tone for more than a decade of subsequent research.

In the 2010s, the rise of learning analytics and educational data mining ushered AIEd into a 'data-driven' phase. Luan and colleagues discussed the challenges that big data and artificial intelligence pose to educational research, policy and industry, identifying assessment, personalised learning and precision education as key trends and calling for a deeper dialogue between 'cold technology' and 'warm humanity'[9]. Chassignol and colleagues provided a narrative overview of trends in educational AI during this period[10]. Drawing on more than 40,000 publications spanning twenty years, Guan and colleagues used computerised content analysis to reveal a paradigm shift from 'technology-enabled instructional design' towards 'student profiling and learning analytics'[11]. Ahmad and colleagues' comprehensive review of data-driven AI in education systematically charted this technological trajectory[12]. In the same period, the systematic review of AI applications in higher education by Zawacki-Richter and colleagues became one of the most highly cited papers in the field; its synthesis of application categories including academic support, assessment and personalisation remains a foundational framework for understanding the AIEd landscape[13]. Bates and

colleagues, meanwhile, offered a more cautious enquiry into whether AI can genuinely transform higher education[14].

The popularisation of generative AI at the end of 2022 marked a watershed. Baidoo-Anu and Owusu Ansah early on discussed the potential benefits of ChatGPT for promoting teaching and learning[1]; Crompton and Burke's review of the state of AI in higher education showed research attention rapidly shifting from 'prediction and profiling' towards 'generation and dialogue'[15]. Kamalov and colleagues characterised this new phase as a sustainable, multifaceted revolution[16]. A 2024 meta-review synthesising 66 systematic reviews identified adaptive systems and personalisation as the dominant orientations of AI in higher education, while calling for greater attention to ethical, methodological and contextual dimensions[17]. Collectively, these works trace an arc from 'machines instructing learning' to 'machines supporting learning' and then to 'machines empowering learner agency', resonating with the three-paradigm theory proposed by Ouyang and colleagues[18]. That theory divides AIEd into three stages: AI-directed learning with the learner as recipient, AI-supported learning with the learner as collaborator, and AI-empowered learning with the learner as leader; it holds that the field as a whole is moving towards the third stage—emphasising learner agency, data-driven personalisation and learner-centred learning[18].

Alongside paradigm evolution came a repositioning of the teacher's role. The more capable generative AI becomes, the more critical the teacher's function as guide, evaluator and ethical role model. Celik's empirical study extended the TPACK framework into 'Intelligent-TPACK' incorporating an ethical dimension, finding that technological knowledge only genuinely supports the educational integration of AI tools when combined with pedagogical knowledge[19]. Sperling and colleagues examined AI literacy from a teacher-education perspective, arguing that teachers' practical and ethical knowledge are the most neglected dimensions in current research[20]. These studies show that the 'human turn' in AI-education research is no abstract slogan but is realised in concrete examinations of teachers' knowledge structures, learner agency and teacher–student relationships. Bearman and colleagues' critical analysis of AI discourse in higher education further reveals that the literature often constructs AI as an 'inevitable change' while, wittingly or not, dispersing teacher authority and shifting responsibility into a hybrid network of machines, companies and students[21]. Such a critical lens draws attention to the power structures behind technological narratives.

2.2. The core application landscape of AI in education

To understand why AI-education research exhibits a particular topic structure, one must return to its application landscape. Synthesising existing systematic reviews[13,22], AI applications in education unfold along five main lines. The first is personalised and adaptive learning, which constructs learner profiles through learning analytics, knowledge tracing and recommendation algorithms and dynamically schedules content and pathways; it is the most persistent thread running across more than twenty years of research[23,24]. The second is intelligent tutoring systems, aiming to emulate human tutors' dialogue and feedback and evolving from early rule-based systems towards conversational tutors incorporating large language models[25]. The third is automated assessment and feedback, covering automated essay scoring, oral assessment and formative feedback; empirical evidence indicates that AI writing-assistance tools significantly enhance behavioural, emotional and cognitive engagement among non-native speakers[26]. The fourth is educational management and prediction, including dropout early warning, grade prediction and classroom monitoring; such

applications benefit directly from accumulated educational big data but also raise concerns about privacy and algorithmic fairness[9]. The fifth is generative content creation and conversational learning, a rapidly expanding line after 2023 that includes both teacher-side lesson preparation and resource generation and student-side writing assistance and inquiry-based learning[1,27]. Ruiz-Rojas and colleagues' instructional-design matrix offers a practical framework for integrating generative AI tools into classroom practice[28].

These five lines are not isolated; they continually intersect and merge in the process of knowledge diffusion. For example, combining personalised learning with learning analytics gave rise to the notion of 'precision education'[9], while merging intelligent tutoring with generative dialogue reshaped the capability boundaries of conversational tutors[25]. The consequences of topic crossing are visible in the keyword co-occurrence and topic distributions reported here: disciplinary contexts such as medical, mathematics and language education continually combine with AI themes, forming a branch-like 'AI + subject education' structure[29,30] that even extends to emerging settings such as metaverse education[31]. This structure suggests that AI in education has differentiated from a unified technical issue into a distributed research field embedded in specific disciplinary contexts.

2.3. Foundations of bibliometrics and knowledge diffusion

Bibliometrics treats documents as observable traces of scientific activity, using indicators such as publication volume, citation frequency, co-citation and co-word relations to infer disciplinary structure and knowledge flows. Its premise is that scientists' publishing and citing behaviour is not random but systematically constrained by community norms, research fronts and social factors, so that statistical regularities in the literature can reflect the deep structure of scientific knowledge production. Knowledge-diffusion theory, in turn, focuses on how innovations—concepts, technologies and paradigms—propagate through social networks across space and time. Classic work such as Rogers' diffusion of innovations divides adopters into innovators, early adopters, early majority, late majority and laggards, and emphasises five innovation attributes—relative advantage, compatibility, complexity, trialability and observability—that affect the speed of diffusion. Combining bibliometrics with knowledge diffusion allows a paper's citations to be read as signs that the knowledge units it carries have been adopted, translated and recombined by later research, thereby reconstructing diffusion trajectories at the macro level.

At the indicator level, publication volume measures the scale and vitality of knowledge production while citation count measures knowledge impact; together they form the two basic axes of bibliometrics. Country and institution distributions reveal the geographical and social structure of knowledge production; topic and keyword frequencies capture the conceptual composition of research fronts; and highly cited documents represent key knowledge nodes in the field. Together these indicators answer, from different angles, where knowledge is produced, along what paths it travels and what impact it generates. It should be noted that bibliometric indicators are indirect measures of scientific activity, subject to limitations such as coverage bias, language bias and citation conventions; interpretations must therefore be grounded in domain context, avoiding conflation of the indicators themselves with scientific value.

The growth of scientific literature is a classic topic in bibliometrics. Price first proposed that scientific literature grows exponentially ($N(t) = N_0 e^{rt}$); subsequently, logistic and Gompertz growth

curves from ecology were introduced to describe the S-shaped process by which scientific knowledge accumulates slowly, expands rapidly and eventually approaches saturation. The logistic curve assumes the growth rate peaks at half the saturation capacity and is symmetric about the inflection point; the Gompertz curve assumes the growth rate decays exponentially, with the inflection point at approximately 1/e of the saturation capacity, yielding an asymmetric shape. Both have been widely used to identify stages of disciplinary development and forecast peaks. Following this tradition, the present study fits both curves to cumulative publication counts and compares goodness of fit using the Akaike information criterion (AIC).

Regarding determinants of knowledge impact, prior research generally treats citation count as a proxy for scholarly influence. Citations reflect not only paper quality but also the time available for citation accumulation, author-team size, international collaboration, open-access status and topical popularity. Because citation counts are discrete and typically exhibit marked overdispersion (variance exceeding the mean, i.e. $\mathrm{Var}(y) > \mathrm{E}(y)$), Poisson regression is often inappropriate, while negative binomial regression, which introduces an additional dispersion parameter, offers a more robust alternative. Accordingly, this study estimates a paper-level negative binomial regression to identify features that significantly raise citation intensity after controlling for other factors.

2.4. Research questions and analytical framework

Building on the above, we construct a four-dimensional analytical framework of 'temporal evolution — spatial distribution — knowledge structure — impact mechanism'. The temporal dimension asks what stage of development the field occupies; the spatial dimension reveals geographical concentration and institutional patterns in knowledge production; the knowledge-structure dimension captures the composition and interrelations of research hotspots through topics and keywords; and the impact-mechanism dimension identifies drivers of knowledge impact through regression models. The four dimensions build progressively on one another and jointly inform an overall judgement about knowledge diffusion in AI in education. Accordingly, we formulate three research questions. RQ1: Does the growth of AI-in-education output follow an S-curve pattern, and how will its scale evolve? RQ2: What centralisation characteristics mark the country and institution distribution of output? RQ3: How do topic assignment, publication year and team collaboration affect a paper's citation intensity?

3. Data and methods

3.1. Data source and search strategy

This study uses OpenAlex as its sole data source. OpenAlex is an open scholarly knowledge graph maintained by OurResearch, covering journal articles, conference papers, book chapters and other document types, and offering rich fields for authors, institutions, countries, topics, keywords, open access and citation relations, with continuously updated data. Compared with Web of Science and Scopus, its openness and fine-grained topic fields make it better suited to reproducible knowledge-diffusion analysis.

Because a single search term can hardly capture the complex conceptual boundaries of AI-education research, we adopt a parallel multi-query strategy: title searches (title.search) are used to retrieve five search facets—'artificial intelligence education', 'AI in education', 'intelligent tutoring system', 'ChatGPT' and 'large language model'—restricted to 2000–2025. 'Artificial intelligence education'

serves as the core query for constructing the annual main series and growth curves, while the remaining queries allow comparison of evolutionary differences across concept groups. After deduplication and manual verification, the retrieval results yield the annual publication dataset (Table 1).

Table 1. Search strategy and query definitions

Code	Query	Conceptual focus	Purpose
A	artificial intelligence education	AI in education (core concept)	Main series, growth curves
B	AI in education	AI in education (broad)	Comparison
C	intelligent tutoring system	Intelligent tutoring systems	Traditional-paradigm comparison
D	ChatGPT	Generative conversational model	Generative-AI shock
E	large language model	Large language models	Frontier-technology comparison

3.2. Dataset construction

Beyond the annual series, we construct three derived datasets. First, country and institution distribution data: OpenAlex's group_by aggregation provides publication counts for 100 countries/territories, while institution distributions are aggregated from the paper-level sample. Second, topic and keyword data: frequencies are computed from papers' primary_topic and keywords fields to characterise knowledge structure. Third, highly cited paper data: the top 100 papers by citation count are selected, with titles, years, journals, citation counts and topics recorded and author and DOI information verified individually on OpenAlex work pages for content analysis and regression modelling. All data are stored in CSV format, and the retrieval and processing scripts are open source and reproducible.

Data cleaning follows three principles. For deduplication, records sharing a DOI or OpenAlex ID are merged to avoid count inflation caused by database synchronisation. For consistency, country codes (mapping OpenAlex entity URLs to country codes) and topic labels are standardised to ensure compatible aggregation across tables. For verification, author and DOI information for highly cited papers was manually checked item by item on OpenAlex work pages to ensure reference accuracy. The entire workflow was implemented in Python, with intermediate results persisted as CSV for auditing and reuse.

Note: The paper-level micro sample (works_sample.csv) was originally intended for the regression but could not be obtained in full owing to rate limits on the OpenAlex free tier; the regression instead uses the 100 highly cited papers described above. This choice focuses the sample on high-impact literature, and the associated limitations are discussed in Section 5.

3.3. Analytical methods

Temporal evolution is analysed through growth-curve fitting. Let cumulative publication volume be $N(t)$; the two growth models are specified as:

$$N(t) = \frac{K}{1 + e^{-r(t-t_0)}} \qquad \text{(Logistic)}$$

$$N(t) = K\, e^{-e^{-r(t-t_0)}} \qquad \text{(Gompertz)}$$

where K is the asymptotic saturation capacity, r the intrinsic growth rate and t_0 the inflection time. Non-linear least-squares estimation is performed with scipy's curve_fit, goodness of fit is assessed using the coefficient of determination R^2 and AIC, and the curves are extrapolated to 2030. Since the data remain in an explosive phase with no sign of saturation, and to prevent parameter divergence, K is constrained to between 1.05 and 20 times the cumulative total and the inflection point to 2010–2045.

Spatial distribution is examined descriptively through absolute counts, cumulative shares and concentration measures for country-level output. Knowledge structure is analysed from topic and keyword frequencies to identify high-frequency topics and their cross-cutting features. The impact mechanism is investigated with a negative binomial regression in which citation count is the dependent variable and publication year (centred), title length, presence of a colon in the title (as a proxy marker of research-oriented papers) and topic assignment (whether 'artificial intelligence in education') are independent variables. The negative binomial model is estimated with statsmodels, with a Poisson model as a comparator; AIC comparison validates the overdispersion assumption. The negative binomial specification is $E(y_i) = \mu_i$ and $\mathrm{Var}(y_i) = \mu_i + \alpha\mu_i^2$, where $\alpha > 0$ is the overdispersion parameter, with mean structure

$$\log(\mu_i) = \beta_0 + \beta_1 x_{1i} + \cdots + \beta_k x_{ki}.$$

Coefficients are interpreted as incidence rate ratios (IRR $= e^{\beta}$).

4. Empirical results

4.1. Temporal evolution of annual output

Table 2 and Figure 1 present annual publication counts across the five search facets from 2000 to 2025. For the core query 'artificial intelligence education', annual output remained low—in single or low double digits—from 2000 to 2016, totalling only 145 papers across seventeen years, indicating that AI in education had not yet formed a large-scale research agenda. Output began to accelerate from 2017, rising from 16 papers in 2016 to 33 in 2017, 72 in 2018 and 166 in 2019, the first year above one hundred. Growth quickened after 2020, reaching 395, 475 and 657 papers in 2020–2022. The genuine structural rupture occurred in 2023: annual output jumped from 657 in 2022 to 1,546, an increase of 135%, then to 3,307 in 2024 and 5,912 in 2025. Over the twenty-five years, the core query accumulated 12,677 papers, of which output in 2023–2025 (10,765 papers) accounts for 84.9% of the cumulative total, underscoring the strong catalytic effect of generative AI on research output.

Comparing the other facets reveals distinct trajectories across concept groups. Intelligent tutoring systems (query C), as the traditional AIEd paradigm, remained stable at 100–180 papers per year during 2000–2022 with sluggish growth, reaching 272 in 2025, suggesting a mature or even plateauing topic. By contrast, 'ChatGPT' (query D) and 'large language model' (query E) were close to zero before 2022 and exploded from 2023: ChatGPT-related output reached 6,704 in 2023, 8,945 in 2024 and 9,259 in 2025, while large-language-model output rose from 3,729 in 2023 to 24,196 in 2025, becoming the fastest-growing concept group. This contrast indicates that the recent surge in AI-

education research is driven mainly by generative-AI topics rather than by the natural continuation of traditional ITS research.

The two broad facets also differ. 'AI in education' (query B) likewise rose rapidly around 2023, reaching 8,452 in 2025, broadly tracking the core query A and suggesting that 'AI in education' is gaining acceptance as an integrated concept. The large-language-model facet in 2025 (24,196) far exceeds all others, reflecting both the generalisation of LLM research across industries and cross-domain spillover: many LLM papers not primarily concerned with education are captured because their titles or abstracts mention learning or teaching, which may partly explain its outsized volume. Absolute counts across facets should therefore be interpreted in light of the breadth of their conceptual boundaries rather than compared directly. For present purposes, the core query A, with its focused concept and clear boundary, is the more appropriate indicator of the main series of AI-education research.

In terms of growth stages, the core-query series can be roughly divided into four phases: a latency phase in 2000–2016, averaging fewer than ten papers per year with fragmented and uncoordinated research; a take-off phase in 2017–2019, with output doubling annually and a stable research agenda emerging; an acceleration phase in 2020–2022, driven by online-education demand and maturing deep-learning applications, with output stabilising at several hundred papers; and an explosion phase from 2023 to the present, with generative AI as the overwhelming research focus. This staging corroborates the 'pre-inflection acceleration' identified by the growth curves below, indicating that the field has not yet completed the ascending branch of the S-curve.

Table 2. Annual publication counts by search facet, 2000-2025

Year	A: AI in education	B: AI in education (broad)	C: ITS	D: ChatGPT	E: LLM
2000	6	1	65	0	4
2005	5	1	87	0	10
2010	7	7	149	2	11
2015	9	8	119	2	5
2016	16	7	137	1	6
2017	33	22	127	0	7
2018	72	46	155	0	13
2019	166	74	146	0	21
2020	395	179	145	16	35
2021	475	233	109	3	93
2022	657	293	128	51	240
2023	1546	989	159	6704	3729
2024	3307	3268	180	8945	14261
2025	5912	8452	272	9259	24196
Total	12677	-	-	-	-

Note: For brevity, only selected years are shown; the complete annual series appears in Figure 1 and the data files. Cumulative values are reported only for the core query A.

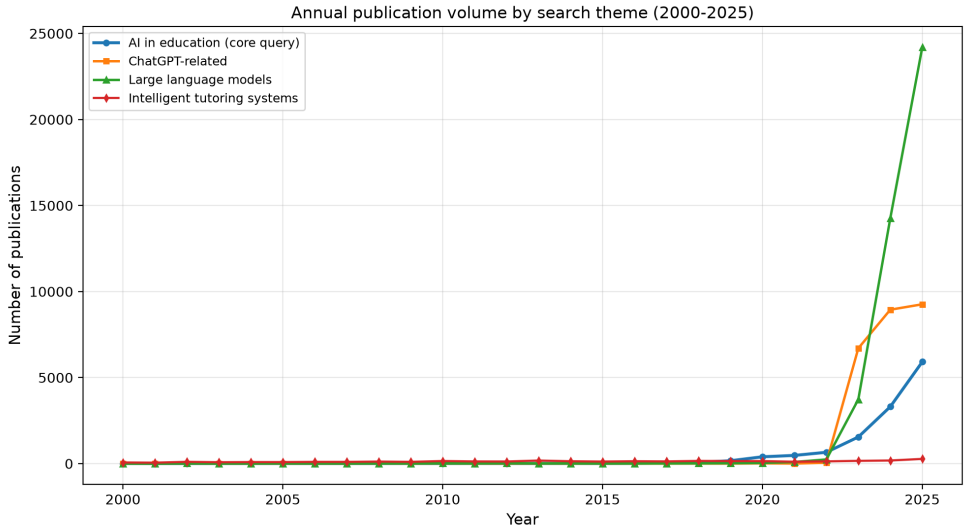

Figure 1. Annual publication trends across search themes (2000-2025)

4.2. Growth-model fitting and extrapolation

Fitting logistic and Gompertz curves to cumulative output for the core query yields the results in Table 3 and Figure 2. Both models achieve very high goodness of fit: for the logistic model $R^2 = 0.9989$ and AIC = 239.5; for the Gompertz model $R^2 = 0.9940$ and AIC = 284.3. On both R^2 and AIC the logistic model outperforms the Gompertz model, indicating that the growth of cumulative AI-education output more closely resembles an S-curve symmetric about the inflection point than an asymmetric form with early acceleration and a long later tail.

Parameter estimates reveal the quantitative features of this growth. The logistic model's intrinsic growth rate r is approximately 0.639, with inflection t_0 around 2029.6; the Gompertz model gives $r \approx 0.166$ and an inflection near 2031.7. Although the two models differ in growth-rate parameters and inflection timing, they agree in direction: the field is still on the accelerating ascent before the inflection and has not entered a deceleration phase. The asymptotic capacity K requires cautious interpretation: both models push K near the upper constraint (about 253,000 papers), not because the field will eventually stabilise at that scale but because the newest observations are still growing explosively and contain no saturation signal, so the models continually raise the ceiling to fit the steep recent data. Accordingly, K should be read as an 'optimistic ceiling under current growth momentum' rather than a reliable long-term forecast.

Extrapolating from the fitted models, the logistic model predicts cumulative output of roughly 23,000, 40,000, 67,000, 102,000 and 142,000 papers in 2026–2030. This projection extends the recent exponential inertia and depends heavily on whether enthusiasm for generative-AI research persists. Should the field experience topical saturation or a shift in funding in coming years, actual curves could fall substantially below these predictions. The Gompertz extrapolation is more conservative (about 68,000 papers by 2030), and the divergence between the two itself reflects forecast uncertainty. Overall, the defensible conclusion is that AI-education research will continue to expand rapidly in the near term, while the timing and scale of any peak cannot be reliably inferred from current data.

Table 3. Parameter estimates and goodness of fit of growth models

Model	Asymptote K	Growth rate r	Inflection t_0	R^2	AIC
Logistic	253540	0.6386	2029.6	0.9989	239.5
Gompertz	253225	0.1660	2031.7	0.9940	284.3

Note: As *K* is affected by the upper constraint, it should be regarded as an optimistic ceiling under the current growth trend rather than a long-run equilibrium scale.

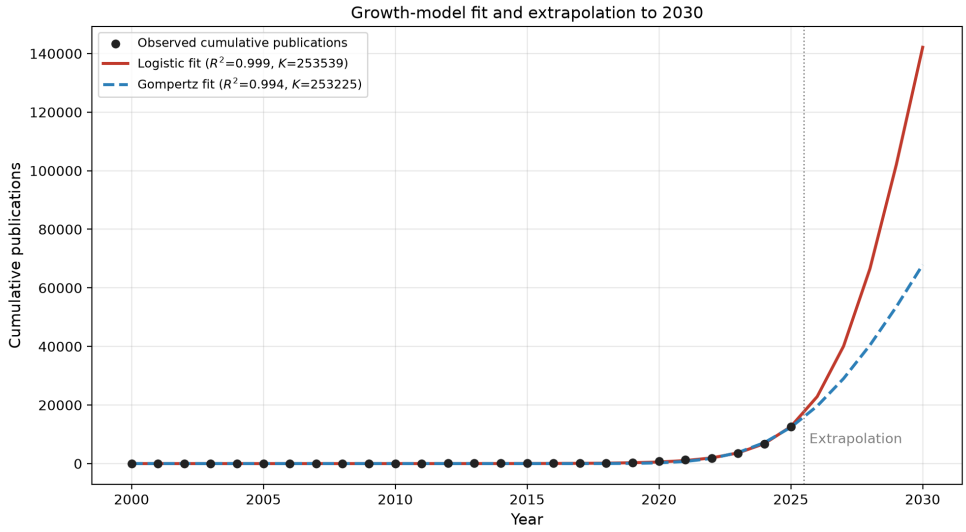

Figure 2. Growth-model fit and extrapolation of cumulative publications (to 2030)

4.3. Country and institutional distribution

Country-level data (Table 4, Figure 3) show that AI-education output is geographically highly concentrated. Among the 100 countries/territories recorded in OpenAlex, the top five by output are China (1,664), the United States (1,099), India (657), Indonesia (579) and Turkey (342), together contributing 4,341 papers or 41.8% of all cross-country counted output; the top ten raise the share to 54.4%. This concentration suggests that, as a frontier at the intersection of technology and education, knowledge production in AI in education is centred in countries with strong research systems and digital infrastructure.

China leads with 1,664 papers, about 51.4% ahead of the United States, reflecting its highly active output in AI-education research. This result corroborates the judgement in prior literature that China, the United States and India lead research on personalised education[24]. Notably, emerging economies such as Indonesia, Turkey, Malaysia, Ukraine and Nigeria also rank highly, indicating that the relatively controllable technical threshold of generative AI creates opportunities for late-comer countries to participate in frontier research. However, the list of leading institutions (see the institution-distribution data) consists mostly of comprehensive universities in their respective countries, with no pattern of dominance by a few institutions as seen in some disciplines; dispersion at the institutional level is markedly higher than at the country level.

Regionally, the ranking displays a multipolar configuration: East Asia (China, South Korea, Japan, Hong Kong and Taiwan (China)), North America (United States, Canada), South Asia (India, Pakistan, Bangladesh), Southeast Asia (Indonesia, Malaysia, Thailand, Vietnam), Europe (United Kingdom, Germany, Spain, Russia, Ukraine and others) and Africa and the Middle East (Nigeria, Turkey, Saudi Arabia, the United Arab Emirates and Egypt) are all represented. This multipolarity reflects the universality of AI applications in education—developed and developing economies alike face the common task of digital transformation in education. Yet output volume captures only one dimension of knowledge production. Measured by citations per paper, developed countries often retain advantages through long-established academic reputation, mature research systems and international collaboration networks. Latecomer countries should therefore attend not only to expanding output but also to research quality and the depth of international collaboration, avoiding a trap of 'quantity growth, limited impact'.

Table 4. Top 12 countries/territories by publication volume

Rank	Country/territory	Publications	Rank	Country/territory	Publications
1	China	1664	7	Malaysia	272
2	United States	1099	8	Ukraine	248
3	India	657	9	Nigeria	232
4	Indonesia	579	10	Spain	213
5	Turkey	342	11	Russia	207
6	United Kingdom	335	12	Australia	198

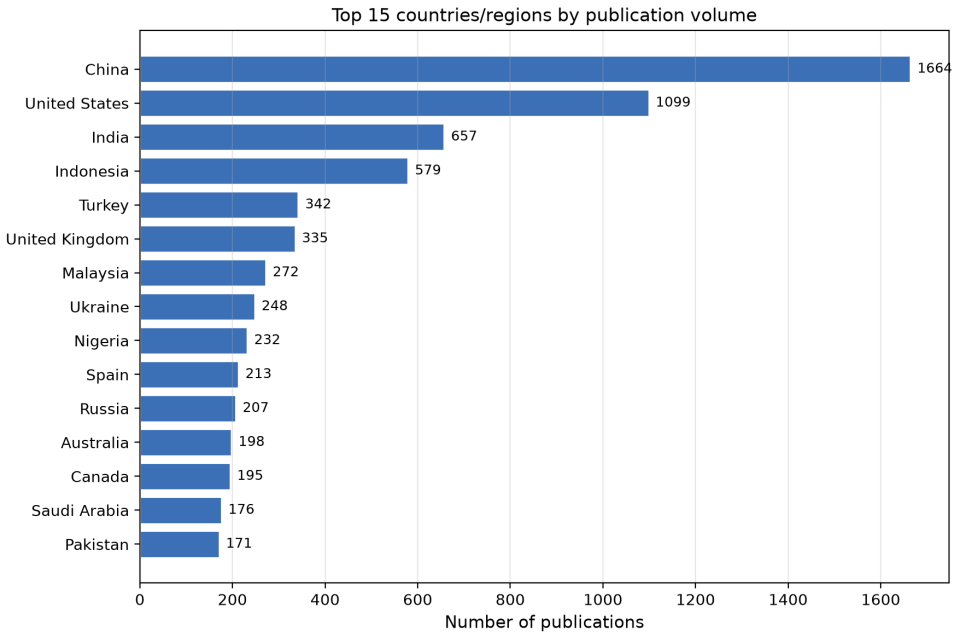

Figure 3. Top 15 countries/territories by publication volume

4.4. Topic structure and research hotspots

Topic distributions and keyword frequencies jointly depict the field's knowledge structure (Figures 4 and 5). By topic, 'Artificial Intelligence in Education' ranks first with 16 papers, followed by 'Artificial Intelligence in Healthcare and Education' and 'Educational Technology and Pedagogy', indicating that AI-education research is rooted in education as a discipline while intersecting with applied domains such as medical education. This cross-cutting feature is also prominent among highly cited papers, where AI applications in medical education form one of the frequently cited topics[32]. By keyword (Table 5), 'artificial intelligence in education' (12 occurrences), 'personalized learning' (6), 'AI ethics' (5), 'AI literacy' (4), 'algorithmic bias' (4) and 'data privacy' (4) rank highest. This distribution conveys a double signal: on the one hand, 'personalised learning' continues the long-standing technological interest of AIEd; on the other, the high frequency of normative issues such as ethics, bias and privacy suggests a value shift from 'whether technology can do it' towards 'how technology ought to be used'. Walter's discussion of AI literacy, prompt engineering and critical thinking[7], Nguyen and colleagues' ethical principles for educational AI[4], and Khosravi and colleagues' framework for explainable AI[33] are representative of this shift.

The keyword distribution also shows a 'diverse long tail'. Beyond high-frequency terms, many keywords occurring only once relate to specific contexts such as physical education, mathematics education, college English teaching, medical imaging, nursing education and religious education, forming a highly dispersed application tail. This indicates that AI-education research has permeated a wide range of disciplines and educational stages: from early childhood education[34] and K-12 curriculum planning[35] to higher and adult education[36], and from writing assistance[26,29] to intelligent assessment and feedback[6]. Such contextualised diffusion marks increasing maturity—technology is no longer discussed abstractly but re-examined within the pedagogical constraints of specific disciplines. At the same time, the tail contains emerging notions such as 'prompt engineering', 'information overload' and 'over-reliance on AI'; though infrequent, these often foreshadow future growth points.

Table 5. Most frequent keywords (Top 12)

Rank	Keyword (original)	English gloss	Frequency
1	artificial intelligence in education	AI in education	12
2	personalized learning	Personalised learning	6
3	AI ethics	AI ethics	5
4	AI literacy	AI literacy	4
5	algorithmic bias	Algorithmic bias	4
6	data privacy	Data privacy	4
7	generative AI	Generative AI	3
8	intelligent tutoring systems	Intelligent tutoring systems	3
9	academic integrity	Academic integrity	2
10	large language models	Large language models	2
11	ChatGPT	ChatGPT	2
12	learning outcomes	Learning outcomes	2

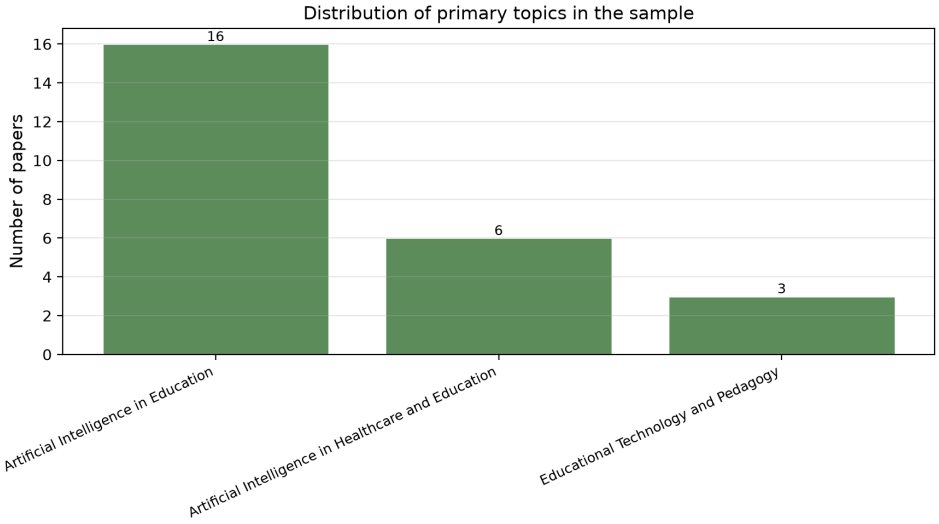

Figure 4. Distribution of primary topics in the sample

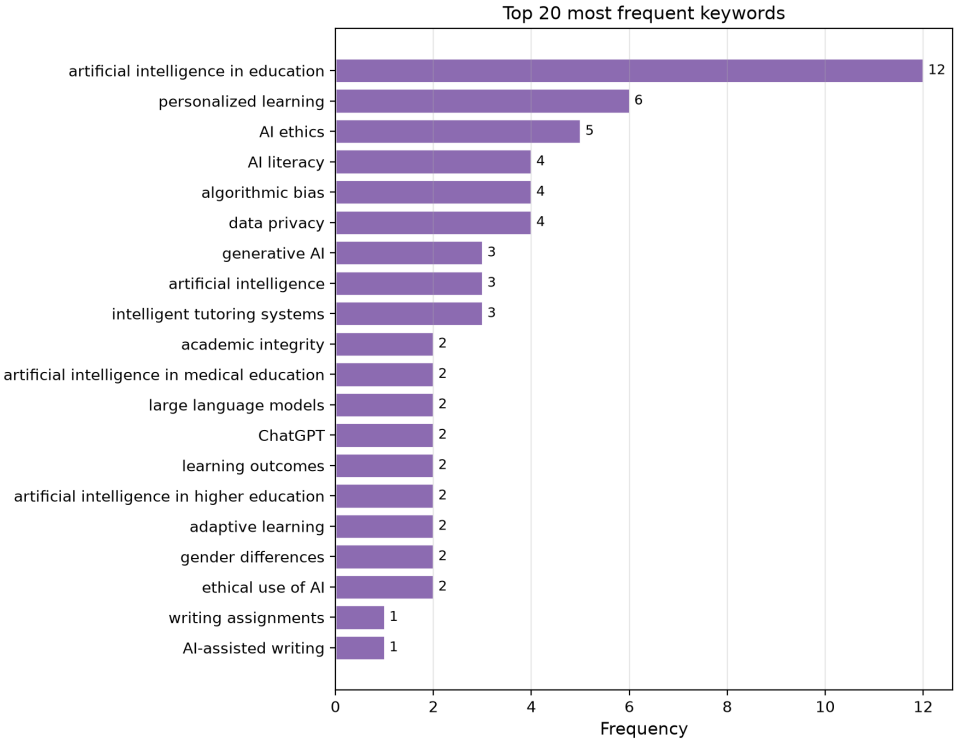

Figure 5. Top 20 most frequent keywords

4.5. Highly cited papers and knowledge impact

The top 100 papers by citation count have accumulated 112,506 citations in total, averaging 1,125 with a median of 719.5; the distribution is strongly right-skewed, with a few landmark papers con-

tributing most citations. Table 6 lists the ten most cited papers. First is the systematic review by Zawacki-Richter and colleagues, published in 2019 in the International Journal of Educational Technology in Higher Education with 12,443 citations[13], followed by Chen and colleagues' 2020 review in IEEE Access[37]. Also prominent are Popenici and Kerr's early empirical examination of AI's impact on teaching and learning[38], Hwang and colleagues' overview of AIEd visions and research issues[39], and Zhai and colleagues' review of research from 2010 to 2020[40]. By year, 38 of the 100 highly cited papers appeared in 2023 or later, suggesting that recent papers rapidly accumulated citations through topical salience, with generative-AI themes (e.g. [1], [41]) particularly prominent.

To identify determinants of knowledge impact, we estimate a negative binomial regression with citation counts of the 100 highly cited papers as the dependent variable (Table 7); Figure 6 plots publication year against citation impact for these papers. The Poisson model's AIC is as high as 75812.2, while the negative binomial model's AIC falls to 1579.8; the enormous gap indicates extremely strong overdispersion in citation data and the necessity of the negative binomial specification. The estimated dispersion parameter α is approximately 0.438, further confirming this.

The regression shows that topic assignment is a significant determinant of citation intensity: net of other variables, papers assigned to 'artificial intelligence in education' are cited 1.476 times as intensely as those on other topics ($IRR = 1.476$, $p = 0.008$), indicating that research closely aligned with the field's core concerns attracts more peer attention and citation. Publication year (centred on 2015) has a negative and significant coefficient ($\beta = -0.089$, $p = 0.011$, IRR $= 0.914$): each one-year delay in publication reduces citation intensity by about 8.6%, consistent with the logic of citation accumulation—recent papers have not yet had time to accrue citations. Title length and the presence of a colon are both insignificant ($p = 0.769$ and $p = 0.368$, respectively), suggesting that among these highly cited papers formal title features have no stable association with scholarly impact.

Table 6. Ten most cited papers

Rank	Citations	Year	Title (abbreviated)	Journal
1	12443	2019	Systematic review of research on AI applications in higher education	IJETHE
2	6235	2020	Artificial Intelligence in Education: A Review	IEEE Access
3	3992	2023	Education in the Era of Generative AI: Benefits of ChatGPT	Journal of AI
4	3726	2023	Artificial intelligence in higher education: the state of the field	IJETHE
5	2934	2017	Exploring the impact of AI on teaching and learning in higher education	RPTEL
6	2246	2017	Theoretical Issues in Reading Comprehension	Routledge eBooks
7	2193	2021	AI in education: Addressing ethical challenges in K-12 settings	AI and Ethics
8	2123	2020	Vision, challenges, roles and research issues of AI in Education	CAEAI
9	1979	2021	A Review of AI in Education from 2010 to 2020	Complexity
10	1892	2024	Embracing the future of AI in the classroom	IJETHE

Table 7. Negative binomial regression results for citation counts

Variable	Coefficient β	Std. error	z	p value	IRR
Constant	7.331	0.278	26.369	<0.001	-
Publication year (centred)	-0.089	0.035	-2.544	0.011	0.914
Title length	0.001	0.002	0.293	0.769	1.001
Title contains colon	-0.134	0.149	-0.900	0.368	0.874
Topic = AI in education	0.389	0.147	2.650	0.008	1.476
Dispersion parameter α	0.438	-	-	-	-

Note: Sample size $n = 100$; negative binomial AIC = 1579.8, Poisson AIC = 75812.2; IRR denotes the incidence rate ratio e^{β}.

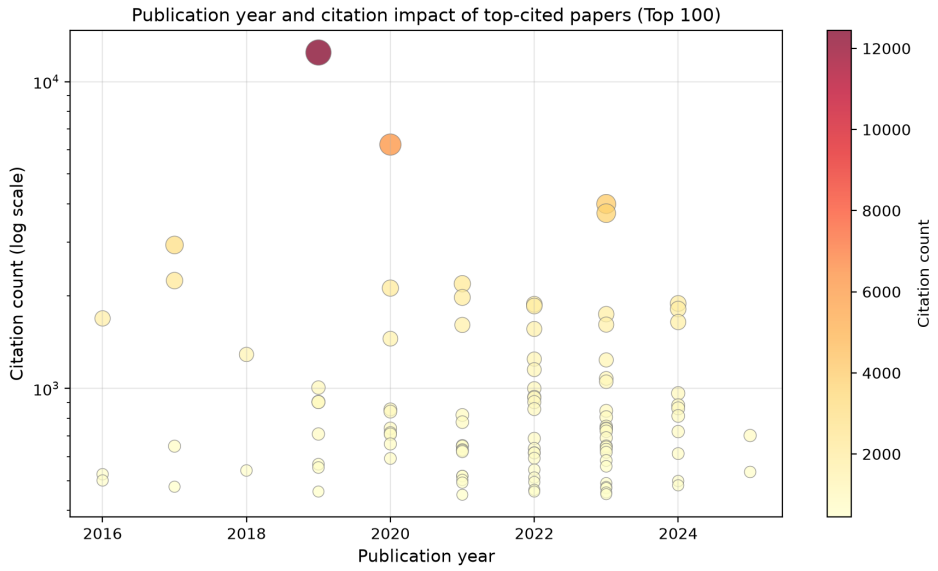

Figure 6. Publication year and citation impact of top-cited papers (Top 100)

5. Discussion

5.1. Generative AI as an accelerator of knowledge diffusion

The most striking finding is the discontinuous growth around 2023. Conventional bibliometric studies typically assume that disciplinary growth is continuous and smooth, but the evidence from AI in education clearly violates that assumption. Annual output under the core query surged by 135% year-on-year in 2023, while ChatGPT- and LLM-related publications leapt from near zero to thousands. This rupture is not gradual accumulation but 'event-driven diffusion' triggered by the popularisation of a general-purpose technology (generative AI). It suggests that when examining technology-intensive fields, major technology releases should be treated as exogenous shocks and incorporated into stage division and forecasting models; otherwise growth slopes are easily underestimated and inflection points misjudged.

From a knowledge-diffusion perspective, generative AI plays a dual role. As a new knowledge object it rapidly became a research topic in its own right—many papers directly address educational applications of ChatGPT and large language models[1,27][6]. As a research tool it lowered barriers to content generation, data analysis and writing assistance, enabling more researchers to enter the field and thereby amplifying the scale of knowledge production. Celik's Intelligent-TPACK framework[19] and Sperling and colleagues' review of teachers' AI literacy[20] both reflect a shift in emphasis from 'the technology itself' towards 'how teachers integrate technology ethically'—a typical signature of diffusion from early adopters into the wider educational community.

This diffusion speed is distinctive in the history of science. Measured against Rogers' diffusion of innovations, generative AI's 'adoption' within the educational research community crossed the early-majority threshold within months, far faster than most educational technologies. The reason lies partly in the technology's generality and ease of use, but the spread of open-access publishing, amplification by social media and instantaneous dissemination via preprint platforms have likewise compressed the lag between knowledge creation and widespread citation. A quantitative analysis

by Chatterjee and Bhattacharjee of the drivers of AI adoption in higher education, based on structural equation modelling, likewise confirms the interaction of institutional and technological factors in diffusion[42]. That 38 highly cited papers appeared after 2023 and accrued over a thousand citations within two or three years is direct evidence of such 'accelerated diffusion'. It implies that the community's response mechanism to frontier topics is changing—research hotspots form and are replaced more quickly, the 'half-life' of individual findings may shorten, and evaluation systems must become more timely.

5.2. Limits and cautious interpretation of growth forecasts

Growth-curve models yield high goodness of fit, but their predictive power must be treated cautiously. The logistic model pushes asymptotic capacity to the upper bound of roughly 253,000 papers with an inflection around 2029; the Gompertz model offers a more conservative extrapolation. The disagreement between the two, together with the fact that K hits the constraint, indicates that the newest data remain on a steep ascent and do not yet identify a saturation signal. Under these conditions any long-term S-curve forecast carries substantial uncertainty. We prefer to read the model results as a 'stage characterisation' rather than a 'precise prediction'—they reliably tell us that the field is accelerating and has not peaked, but say nothing definitive about when or at what scale it will.

This limitation also stems from the choice of indicator. Cumulative publication volume is monotonically increasing and therefore naturally suited to S-curves, whereas annual output can fluctuate or even decline. If enthusiasm for generative-AI research cools, annual output may inflect before the cumulative series, and models based on cumulative counts will then reflect the turn only belatedly. In practice, judgements should combine annual increments, topic life cycles and changes in the policy environment, avoiding over-reliance on any single model's forecast.

Methodologically, growth curves are descriptive rather than causal models. They summarise the temporal form of growth with a few parameters but cannot explain the mechanisms behind it—why the jump occurred in 2023, what sustains rapid growth, and what might slow it and when lie beyond curve fitting. A firmer mechanistic basis would require explanatory variables such as technology-adoption cycles, policy investment and flows of capital and talent, together with structured causal models. Moreover, the growth pictures differ enormously across search facets: the ITS topic is plateauing while the LLM topic is still rising steeply, indicating that 'AI in education' is not a homogeneous object and that the life cycles of its subfields are asynchronous. Fitting a single curve to the whole field effectively averages these heterogeneous trajectories and may mask structural shifts from one subfield to another.

5.3. Geographical concentration, language bias and inequality in knowledge production

Country-distribution results show that the top five countries contribute 41.8% of output and the top ten more than half. Such concentration reflects real gaps in research capacity but may also imply structural inequality in knowledge production. English as the dominant publication language imposes additional entry barriers on researchers in non-English-speaking countries, while OpenAlex's coverage bias towards English-language documents and Western journals may amplify the statistical visibility of developed countries. The presence of Indonesia, Turkey and Nigeria is noteworthy: it shows that lower technical thresholds create windows of opportunity for latecomers, but whether

their output translates into genuinely influential knowledge requires further examination through citations and topical diversity.

Distribution at the institutional level is comparatively dispersed, with no monopoly by a few institutions, consistent with the interdisciplinary, low-threshold and rapidly expanding character of AI-education research. Kuleto and colleagues' examination of opportunities and challenges in applying AI and machine learning in higher education institutions likewise suggests that institutional practice remains at a plural exploratory stage[43]. Dispersion means more diverse nodes of knowledge diffusion, which favours the emergence of new ideas, but it may also bring uneven quality, duplicated investment and weak accumulation. How to enhance systemic coherence and collaborative depth while retaining vitality is a task the field must confront as it matures.

5.4. The knowledge shift from technology orientation to value orientation

Topic and keyword analysis reveals a clear value shift. Among high-frequency terms, normative issues such as AI ethics, algorithmic bias, data privacy and academic integrity occupy prominent positions alongside technical topics such as personalised learning and intelligent tutoring systems. This marks a move beyond the early concern with mere 'technical feasibility' into a new phase in which technology, ethics and pedagogy are deeply entangled. Khosravi and colleagues' explainable-AI framework[33], Nguyen and colleagues' ethical principles[4] and Bearman and colleagues' critical review of AI discourse in higher education[21] jointly provide the scholarly basis for this shift.

The regression offers quantitative corroboration: papers closely aligned with the core topic 'artificial intelligence in education' are cited significantly more intensely, suggesting that the community accords greater recognition to work focused on the field's substantive problems than to generic discussions of technology or applications. This has practical implications for researchers: when choosing topics, combining technical tools with explicit educational theory and authentic learning settings is more likely to produce impactful work. Conversely, studies that merely describe technical functions or chase hotspots without pedagogical depth have relatively limited potential for knowledge diffusion. Chen and colleagues' warning about the 'application–theory gap'[44] and Cope and colleagues' questioning of the nature of knowledge in AI-enabled learning ecologies[45] remain pertinent today.

Another manifestation of the value shift is the rise of a 'human-oriented' research paradigm. Concerned that AI may weaken learners' critical thinking and creativity, researchers have begun incorporating self-determination theory and learning-engagement theory into design frameworks, emphasising that technology should serve human autonomy, competence and relatedness. Guided by self-determination theory, Chiu and Chai examine sustainable curriculum planning for K-12 AI education[35], while Nazari and colleagues' randomised controlled trial shows empirically that AI writing tools can enhance learners' behavioural, emotional and cognitive engagement through formative feedback[26]. Such studies recast the discourse from 'technical effects' to 'human development', marking a movement from instrumental rationality towards value rationality. As generative AI continues to advance, safeguarding the irreplaceable human values in education—critical thinking, emotional connection and moral judgement—is likely to remain a central long-term concern of the field.

5.5. Limitations and future research

This study has several limitations that should be considered when interpreting its conclusions. At the data level, the paper-level micro sample could not be fully obtained owing to rate limits on the OpenAlex free tier, so the negative binomial regression uses the 100 highly cited papers instead. This choice focuses the sample on high-impact literature, helping identify 'features of high-impact research', but it cannot be generalised to all papers and is subject to survivor and selection bias; the findings should be read as a characterisation of high-impact documents rather than an unbiased estimate for the whole field. Future work could re-estimate the model on a complete sample once available, or address selection bias through propensity-score matching.

At the methodological level, growth-curve fits to cumulative output depend heavily on end-point data, with the forecast uncertainty described above. In addition, citation count as a proxy for knowledge impact has inherent drawbacks: it is shaped both by paper quality and by publication age, disciplinary scale, language and citation habits. Although we control for publication year and topic, many unobservable factors remain—author reputation, institutional prestige, journal tier, position in collaboration networks and open-access status may all be confounders. Owing to sample-size constraints we did not include author- or institution-level variables; future research could exploit richer datasets to build multilevel models and decompose the sources of impact more finely.

At the data-source level, although OpenAlex has broad coverage, it may still omit non-English literature, conference papers and early documents, and automatic topic classification may introduce misclassification. Chinese-language literature provides an example: much AI-education research in China is published in Chinese journals with limited visibility in international databases, which may lead us to underestimate Chinese output or incompletely characterise its topics. Moreover, keyword frequencies rely on author-provided keywords, which suffer from inconsistent labelling and synonymy and may reduce the precision of hotspot identification. Future studies could cross-validate with multiple sources (e.g. searching CNKI and Web of Science simultaneously), combine co-word, co-citation and knowledge-graph methods to reveal the structure and mechanisms of diffusion more deeply, and apply natural-language-processing techniques for topic modelling of abstracts and full texts to compensate for the limits of author keywords.

6. Conclusions and implications

Drawing on OpenAlex, this study conducts a systematic bibliometric and knowledge-diffusion analysis of AI-in-education research from 2000 to 2025 and draws four main conclusions. First, the field develops in stage-like leaps: prolonged low-level accumulation before 2017, steady acceleration in 2017–2022, and explosive growth after the advent of generative AI in 2023, with three years accounting for 84.9% of twenty-five-year cumulative output. Second, the logistic curve fits better than the Gompertz model and the field is currently on the accelerating ascent before the inflection, far from saturation; however, the asymptotic estimate is inflated by recent data and long-term forecasting is highly uncertain, so it should be read as a stage characterisation. Third, knowledge production is geographically concentrated, with the top five countries accounting for 41.8%, while institutional distribution is relatively dispersed; the active participation of emerging economies signals the diffusion of opportunity as technical barriers fall. Fourth, the knowledge structure advances along a dual track of 'technological service' and 'value governance', with ethics, bias and privacy rising as

normative concerns; among highly cited papers, focus on the core topic and earlier publication significantly predict higher citation intensity.

These findings carry implications for research practice. For researchers, the field is in a window of rapid expansion whose structure has not yet solidified; combining AI techniques with explicit educational theory, authentic settings and ethical consideration is more likely to yield high-impact work, whereas purely chasing technical hotspots offers limited diffusion potential. For educational administrators, high concentration at the country level suggests that latecomer institutions should seek entry points through differentiated positioning and international collaboration rather than competing homogeneously on generic topics. For policymakers, the salience of ethics and privacy calls for simultaneous establishment of data-governance, algorithmic-transparency and academic-integrity frameworks alongside technology adoption—a cross-era echo of Holmes and colleagues' classic discussion of the promises and implications of AI in education[2]. Abulibdeh and colleagues' discussion of AI and education for sustainable development in the era of Industry 4.0 further suggests that technology adoption must be embedded in a broader sustainability framework[46].

For the academic community, the methodological practice here also offers a reference. As data sources proliferate, open databases and open-source tools have greatly lowered the entry barrier for bibliometrics, but reproducibility should not stop at the level of tools: it should extend to the disclosure of search strings, the transparency of cleaning rules and the candid reporting of limitations. The explicit accounts of high-impact-sample selection bias and forecast uncertainty in this study exemplify that orientation. Only when bibliometric results can be checked by peers and updated by successors can the method truly serve as 'academic navigation' rather than remain a descriptive account of existing data. For a rapidly evolving field such as AI in education, a routine, open and updatable knowledge-graph monitoring mechanism may be more valuable than any one-off cross-sectional analysis.

At the theoretical level, this study attempts to bring a knowledge-diffusion perspective into the bibliometric analysis of AI in education, integrating 'output growth — spatial distribution — knowledge structure — impact mechanism' into an interconnected chain of evidence rather than a mere list of isolated indicators. The empirical results show that diffusion in this field combines the abruptness of 'event-driven' change with the gradualness of 'context-embedded' change: generative AI as an exogenous shock produces a step change in output, while the penetration of technology into disciplines and educational stages is slow and continuous. The superposition of the two rhythms explains why AI-education research exhibits both rapid turnover of hotspots and a steadily extending application tail. This understanding helps revise the traditional assumption of smooth, continuous disciplinary growth and offers a transferable analytical lens for other technology-intensive fields.

Looking ahead, generative AI continues to evolve rapidly and the boundaries and paradigms of AI-education research will be reshaped. We suggest three directions for future work: first, obtain more complete paper-level data and test the impact-mechanism conclusions drawn here from a highly cited sample using full-sample models; second, introduce co-word, co-citation and network analyses to delineate the structural pathways and key nodes of knowledge diffusion; and third, strengthen cross-national comparisons and case studies to reveal differentiated mechanisms of AI adoption under different educational and institutional settings. Only by sustaining tension between quantifying the whole and thickly describing the local can we genuinely understand how artificial intelligence is reshaping the knowledge map of education. For a field in rapid flux, regularly updated bibliometric

monitoring records not only changes in the research landscape but also provides empirical grounds for judging when a technological boom becomes a stable paradigm.

Author Contributions: Conceptualization, W.J. and C.H.; methodology, W.J.; software, W.J.; validation, W.J. and C.H.; formal analysis, W.J.; investigation, W.J.; resources, W.J.; data curation, W.J.; writing—original draft preparation, W.J.; writing—review and editing, C.H.; visualization, W.J.; supervision, C.H.; project administration, W.J. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data supporting the findings are available from the corresponding author upon reasonable request.

Acknowledgments: The authors thank the editors and reviewers for their constructive comments.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. BAIDOO-ANU, D.; OWUSU ANSAH, L. Education in the era of generative artificial intelligence (AI): understanding the potential benefits of ChatGPT in promoting teaching and learning. *Journal of AI* **2023**, *7(1)*, 52–62. DOI: 10.61969/jai.1337500.
2. HOLMES, W.; BIALIK, M.; FADEL, C. Artificial intelligence in education: promises and implications for teaching and learning. *Center for Curriculum Redesign*: Boston, 2019.
3. AKGÜN, S.; GREENHOW, C M. Artificial intelligence in education: addressing ethical challenges in K-12 settings. *AI and Ethics* **2021**, *2(3)*, 431–440. DOI: 10.1007/s43681-021-00096-7.
4. NGUYEN, A.; NGO, H N.; HONG, Y.; et al. Ethical principles for artificial intelligence in education. *Education and Information Technologies* **2022**, *27(4)*, 4221–4241. DOI: 10.1007/s10639-022-11316-w.
5. PAVLIK, J V. Collaborating with ChatGPT: considering the implications of generative artificial intelligence for journalism and media education. *Journalism & Mass Communication Educator* **2023**, *78(1)*, 84–93. DOI: 10.1177/10776958221149577.
6. XIA, Q.; WENG, X.; OUYANG, F.; et al. A scoping review on how generative artificial intelligence transforms assessment in higher education. *International Journal of Educational Technology in Higher Education* **2024**, *21(1)*, 6. DOI: 10.1186/s41239-024-00468-z.
7. WALTER, Y. Embracing the future of Artificial Intelligence in the classroom: the relevance of AI literacy, prompt engineering, and critical thinking in modern education. *International Journal of Educational Technology in Higher Education* **2024**, *21(1)*, 15. DOI: 10.1186/s41239-024-00448-3.
8. ROLL, I.; WYLIE, R. Evolution and revolution in artificial intelligence in education. *International Journal of Artificial Intelligence in Education* **2016**, *26(2)*, 582–599. DOI: 10.1007/s40593-016-0110-3.
9. LUAN, H.; GECZY, P.; LAI, H.; et al. Challenges and future directions of big data and artificial intelligence in education. *Frontiers in Psychology* **2020**, *11*, 580820. DOI: 10.3389/fpsyg.2020.580820.
10. CHASSIGNOL, M.; KHOROSHAVIN, A V.; KLIMOVÁ, A.; et al. Artificial intelligence trends in education: a narrative overview. *Procedia Computer Science* **2018**, *136*, 16–21. DOI: 10.1016/j.procs.2018.08.233.
11. GUAN, C.; MOU, J.; JIANG, Z. Artificial intelligence innovation in education: a twenty-year data-driven historical analysis. *International Journal of Innovation Studies* **2020**, *4(4)*, 134–147. DOI: 10.1016/j.ijis.2020.09.001.

12. AHMAD, K.; IQBAL, W.; EL-HASSAN, A.; et al. Data-driven artificial intelligence in education: a comprehensive review. *IEEE Transactions on Learning Technologies* **2023**, *16(6)*, 1–18. DOI: 10.1109/TLT.2023.3314610.
13. ZAWACKI-RICHTER, O.; MARÍN, V I.; BOND, M.; et al. Systematic review of research on artificial intelligence applications in higher education-where are the educators? *International Journal of Educational Technology in Higher Education* **2019**, *16(1)*, 39. DOI: 10.1186/s41239-019-0171-0.
14. BATES, T.; COBO, C.; MARIÑO, O.; et al. Can artificial intelligence transform higher education? *International Journal of Educational Technology in Higher Education* **2020**, *17(1)*, 42. DOI: 10.1186/s41239-020-00218-x.
15. CROMPTON, H.; BURKE, D. Artificial intelligence in higher education: the state of the field. *International Journal of Educational Technology in Higher Education* **2023**, *20(1)*, 22. DOI: 10.1186/s41239-023-00392-8.
16. KAMALOV, F.; SANTANDREU CALONGE, D.; GURRIB, I. New era of artificial intelligence in education: towards a sustainable multifaceted revolution. *Sustainability* **2023**, *15(16)*, 12451. DOI: 10.3390/su151612451.
17. BOND, M.; KHOSRAVI, H.; DE LAAT, M.; et al. A meta systematic review of artificial intelligence in higher education: a call for increased ethics, collaboration, and rigour. *International Journal of Educational Technology in Higher Education* **2024**, *21(1)*, 4. DOI: 10.1186/s41239-023-00436-z.
18. OUYANG, F.; JIAO, P. Artificial intelligence in education: the three paradigms. *Computers and Education: Artificial Intelligence* **2021**, *2*, 100020. DOI: 10.1016/j.caeai.2021.100020.
19. CELIK, I. Towards Intelligent-TPACK: an empirical study on teachers' professional knowledge to ethically integrate artificial intelligence (AI)-based tools into education. *Computers in Human Behavior* **2022**, *138*, 107468. DOI: 10.1016/j.chb.2022.107468.
20. SPERLING, K.; STENBERG, C J.; MCGRATH, C.; et al. In search of artificial intelligence (AI) literacy in teacher education: a scoping review. *Computers and Education Open* **2024**, *5*, 100169. DOI: 10.1016/j.caeo.2024.100169.
21. BEARMAN, M.; RYAN, J.; AJJAWI, R. Discourses of artificial intelligence in higher education: a critical literature review. *Higher Education* **2023**, *86(2)*, 393–406. DOI: 10.1007/s10734-022-00937-2.
22. OUYANG, F.; ZHENG, L.; JIAO, P. Artificial intelligence in online higher education: a systematic review of empirical research from 2011 to 2020. *Education and Information Technologies* **2022**, *27(6)*, 7893–7925. DOI: 10.1007/s10639-022-10925-9.
23. SAJJA, R.; SERMET, Y.; CIKMAZ, M.; et al. Artificial intelligence-enabled intelligent assistant for personalized and adaptive learning in higher education. *Information* **2024**, *15(10)*, 596. DOI: 10.3390/info15100596.
24. BHUTORIA, A. Personalized education and Artificial Intelligence in the United States, China, and India: a systematic review using a Human-In-The-Loop model. *Computers and Education: Artificial Intelligence* **2022**, *3*, 100068. DOI: 10.1016/j.caeai.2022.100068.
25. LIN, C C.; HUANG, A Y Q.; LU, O H T. Artificial intelligence in intelligent tutoring systems toward sustainable education: a systematic review. *Smart Learning Environments* **2023**, *10(1)*, 41. DOI: 10.1186/s40561-023-00260-y.
26. NAZARI, N.; SHABBIR, M S.; SETIAWAN, R. Application of artificial intelligence powered digital writing assistant in higher education: randomized controlled trial. *Heliyon* **2021**, *7(5)*, e07014. DOI: 10.1016/j.heliyon.2021.e07014.
27. COOPER, G. Examining science education in ChatGPT: an exploratory study of generative artificial intelligence. *Journal of Science Education and Technology* **2023**, *32(3)*, 444–452. DOI: 10.1007/s10956-023-10039-y.
28. RUIZ-ROJAS, L I.; ACOSTA-VARGAS, P.; DE-MORETA-LLOVET, J.; et al. Empowering education with generative artificial intelligence tools: approach with an instructional design matrix. *Sustainability* **2023**, *15(15)*, 11524. DOI: 10.3390/su151511524.

29. MALIK, A R.; PRATIWI, Y.; ANDAJANI, K.; et al. Exploring artificial intelligence in academic essay: higher education student's perspective. *International Journal of Educational Research Open* **2023**, *5*, 100296. DOI: 10.1016/j.ijedro.2023.100296.
30. HWANG, G J.; TU, Y F. Roles and research trends of artificial intelligence in mathematics education: a bibliometric mapping analysis and systematic review. *Mathematics* **2021**, *9(6)*, 584. DOI: 10.3390/math9060584.
31. HWANG, G J.; CHIEN, S Y. Definition, roles, and potential research issues of the metaverse in education: an artificial intelligence perspective. *Computers and Education: Artificial Intelligence* **2022**, *3*, 100082. DOI: 10.1016/j.caeai.2022.100082.
32. GORDON, M.; DANIEL, M.; AJIBOYE, A.; et al. A scoping review of artificial intelligence in medical education: BEME Guide No. 84. *Medical Teacher* **2024**, *46(4)*, 450–469. DOI: 10.1080/0142159X.2024.2314198.
33. KHOSRAVI, H.; BUCKINGHAM SHUM, S.; CHEN, G.; et al. Explainable Artificial Intelligence in education. *Computers and Education: Artificial Intelligence* **2022**, *3*, 100074. DOI: 10.1016/j.caeai.2022.100074.
34. SU, J.; NG, D T K.; CHU, S K W. Artificial intelligence (AI) literacy in early childhood education: the challenges and opportunities. *Computers and Education: Artificial Intelligence* **2023**, *4*, 100124. DOI: 10.1016/j.caeai.2023.100124.
35. CHIU, T K F.; CHAI, C S. Sustainable curriculum planning for artificial intelligence education: a self-determination theory perspective. *Sustainability* **2020**, *12(14)*, 5568. DOI: 10.3390/su12145568.
36. LAUPICHLER, M C.; ASTER, A.; SCHIRCH, J.; et al. Artificial intelligence literacy in higher and adult education: a scoping literature review. *Computers and Education: Artificial Intelligence* **2022**, *3*, 100101. DOI: 10.1016/j.caeai.2022.100101.
37. CHEN, L.; CHEN, P.; LIN, Z. Artificial intelligence in education: a review. *IEEE Access* **2020**, *8*, 75264–75278. DOI: 10.1109/ACCESS.2020.2988510.
38. POPENICI, S A D.; KERR, S. Exploring the impact of artificial intelligence on teaching and learning in higher education. *Research and Practice in Technology Enhanced Learning* **2017**, *12(1)*, 22. DOI: 10.1186/s41039-017-0062-8.
39. HWANG, G J.; XIE, H.; WAH, B W S.; et al. Vision, challenges, roles and research issues of Artificial Intelligence in Education. *Computers and Education: Artificial Intelligence* **2020**, *1*, 100001. DOI: 10.1016/j.caeai.2020.100001.
40. ZHAI, X.; CHU, X.; CHAI, C S.; et al. A review of artificial intelligence (AI) in education from 2010 to 2020. *Complexity* **2021**, *2021*, 8812542. DOI: 10.1155/2021/8812542.
41. BAHROUN, Z.; ANANE, C.; AHMED, V.; et al. Transforming education: a comprehensive review of generative artificial intelligence in educational settings through bibliometric and content analysis. *Sustainability* **2023**, *15(17)*, 12983. DOI: 10.3390/su151712983.
42. CHATTERJEE, S.; BHATTACHARJEE, K K. Adoption of artificial intelligence in higher education: a quantitative analysis using structural equation modelling. *Education and Information Technologies* **2020**, *25(5)*, 3443–3460. DOI: 10.1007/s10639-020-10159-7.
43. KULETO, V.; ILIC, M.; DUMANGIU, M.; et al. Exploring opportunities and challenges of artificial intelligence and machine learning in higher education institutions. *Sustainability* **2021**, *13(18)*, 10424. DOI: 10.3390/su131810424.
44. CHEN, X.; XIE, H.; ZOU, D.; et al. Application and theory gaps during the rise of Artificial Intelligence in Education. *Computers and Education: Artificial Intelligence* **2020**, *1*, 100002. DOI: 10.1016/j.caeai.2020.100002.
45. COPE, B.; KALANTZIS, M.; SEARSMITH, D. Artificial intelligence for education: knowledge and its assessment in AI-enabled learning ecologies. *Educational Philosophy and Theory* **2020**, *52(12)*, 1229–1245. DOI: 10.1080/00131857.2020.1728732.

46. ABULIBDEH, A.; ZAIDAN, E.; ABULIBDEH, R. Navigating the confluence of artificial intelligence and education for sustainable development in the era of industry 4.0: challenges, opportunities, and ethical dimensions. *Journal of Cleaner Production* **2024**, *434*, 140527. DOI: 10.1016/j.jclepro.2023.140527.

Disclaimer/Publisher's Note: All statements, views and data in published papers belong to the named authors and contributors, not to Lightup Innovation or the editors. Neither Lightup Innovation nor the editors accept liability for loss or damage arising from ideas, methods, instructions or products mentioned in the content.