**BSIJ** *Business System & Innovation Journal*
Research • Innovation • Digital Transformation

YAYASAN FATHUR RAHMAN
Vol. 01, Issue 03, 2026 | e-ISSN 3163-2831

# From Compliance to Responsible AI: A Bibliometric Study on Educational Data Governance and Student Privacy (2015-2025)

**Ying Chen[1], Ruofan Lin[1], Weijian Kong[1], Yundong Wu[1]*, Shun Guo[1], Jiani Wu[1], Tingting Lv[1], Yuxuan Chen[1]**

**OPEN ACCESS**
*No article processing charge (APC)*

**ARTICLE INFORMATION**

**Page**
142-158

**Received**
28 September 2026

**Revised**
13 October 2026
25 October 2026

**Accepted**
26 October 2026

**Published**
30 October 2026

**DOI**
https://doi.org/10.67955/bsij.v1i3.29

**Copyright**
© 2026 by the author(s). Articles published in this journal are licensed under the Creative Commons Attribution 4.0 International (CC BY 4.0) license, which permits others to copy, distribute, display, and adapt the work, provided that appropriate credit is given to the original authors and to the journal as the first place of publication.

**1** Higher Education, Psychology, and Policy Analysis Laboratory, Naju, 58250, South Korea
* Correspondence: ethanwu0525@gmail.com

**Abstract.** This study maps the development of research on educational data governance and student privacy between 2015 and 2025. It follows the structure of a bibliometric trend review and uses a Scopus-derived record set of 313 articles to analyze the analytical workflow, including annual output, keyword co-occurrence, thematic evolution, policy turning points, and research collaboration. The study is guided by three questions: how the volume of research changed, how the dominant vocabulary moved from learning analytics and ownership to compliance and responsible artificial intelligence, and which problems remained unresolved across the period. The results show five connected phases. From 2015 to 2017, research was organized around learning analytics, educational data mining, student tracking, data ownership, and informed consent. After the General Data Protection Regulation became applicable on May 25, 2018, the literature moved toward data protection, compliance, governance frameworks, and institutional responsibility. The COVID-19 period increased attention to remote learning privacy, online proctoring, surveillance, and vendor data sharing. From late 2022, generative artificial intelligence, ChatGPT, algorithmic bias, ethical artificial intelligence, and artificial intelligence literacy became central. The 2024-2025 frontier connected responsible artificial intelligence with high-risk education systems, human oversight, auditability, and data minimization in large language models. The study argues that compliance is necessary but insufficient. A mature governance model must connect legal duties, institutional processes, technical controls, and educational judgment.

**Keywords:** educational data governance; student privacy; learning analytics; GDPR; online proctoring;

## 1. Introduction

Educational institutions increasingly depend on data systems to support teaching, student services, assessment, quality assurance, and institutional planning. Learning management systems record clicks, submissions, grades, attendance, and communication patterns. Student information systems connect these records with demographic and administrative information. Cloud platforms add another layer because many educational services are provided by external vendors. These systems create opportunities for early support and better resource allocation, but they also change the meaning of privacy in education. Privacy is no longer limited to whether a file is kept secret. It also includes what data are collected, how long they are kept, how they are combined, who can infer new information, and how an automated decision may affect a student (Crawford & Boyd, 2012).

The expansion of learning analytics made this problem visible in the middle of the 2010s (Ferguson, 2012). Early discussions asked whether institutions could use traces of student activity to improve retention and achievement while still respecting student autonomy and informed consent (Siemens & Long, 2011). The debate soon moved to questions of data ownership, institutional power, transparency, and the unequal ability of students to refuse data collection (Slade & Prinsloo, 2013). This was also a problem of contextual integrity because information appropriate in one educational setting could become inappropriate when transferred to another (Nissenbaum, 2004).

The legal environment changed the vocabulary of the field. The General Data Protection Regulation became applicable in the European Union on May 25, 2018, and placed stronger emphasis on lawfulness, purpose limitation, data minimization, rights of data subjects, and organizational accountability (European Parliament and Council of the European Union, 2016). This change did not create the idea of educational privacy, but it gave institutions a clearer legal language for discussing it. Research began to connect learning analytics with compliance, data protection officers, records of processing, impact assessment, and cross-institutional data sharing.

The COVID-19 pandemic then changed the practical context of data governance. During the rapid move to online learning, institutions adopted remote teaching platforms and online proctoring tools under time pressure. Camera access, identity checks, browser monitoring, biometric inference, and the transfer of student data to vendors became normal in

some settings. The central question was no longer only whether a university could collect data for learning support. It was whether emergency surveillance could be justified, limited, and reviewed when students had little choice about the digital environment in which education occurred (Williamson & Eynon, 2020).

Since late 2022, generative artificial intelligence has added a further shift. Large language models can process student work, generate feedback, support tutoring, and assist administrative prediction (Zawacki-Richter et al., 2019). At the same time, they can reproduce social bias, expose personal information through prompts or training data, and make decisions difficult to explain (Bender et al., 2021). The focus of governance has therefore widened from protection of stored records to the management of dynamic systems that infer, rank, recommend, and generate.

This study examines that movement through a bibliometric lens. The purpose is not to claim that publication growth automatically represents better practice. Publication counts show attention, while keywords and thematic links show how a field describes its concerns. A bibliometric review is useful because it can connect changes in vocabulary with policy and technology events across a long period (Zupic & Cater, 2015). The study uses five periods that reflect the main transitions identified in the literature: 2015-2017, 2018-2019, 2020-2021, 2022-2023, and 2024-2025. The last period also includes a policy frontier related to the European Union Artificial Intelligence Act, which entered into force on August 1, 2024 and applies through a phased timetable rather than through one single date (European Parliament and Council of the European Union, 2024).

Three research questions guide the analysis. First, how did publication activity on educational data governance and student privacy change from 2015 to 2025? Second, how did the field's main concepts move from learning analytics and ownership toward compliance, surveillance, algorithmic fairness, and responsible AI? Third, what governance gaps remain when privacy, educational value, institutional accountability, and automated decision making are considered together?

The study maps output and themes, connects them with policy and technology contexts, and proposes privacy governance as a continuing institutional responsibility.

## 2. Research Scope and Method

### *2.1. Search scope*

The search period for the research workflow was August-September 2026, with the Scopus search conducted on September 18, 2026 and the final export downloaded on the same date.

The target corpus is defined by title, abstract, and author keywords. The document type is limited to research articles, while reviews and conference papers are excluded from the main count to reduce variation in publication practices. English-language records are retained for the main analysis because the coding framework and citation audit are conducted in English. Records are expected to cover higher education, schools, educational technology providers, or public education systems. Papers that discuss privacy only in health, employment, or consumer contexts are excluded unless education is a clear setting.

The following query is used in the manuscript as the reproducible search specification. It is written in Scopus advanced search syntax.

TITLE-ABS-KEY ( ("learning analytics" OR "educational data mining" OR "educational data" OR "student data" OR "student information system*" OR "learning management system*") AND (privacy OR "data governance" OR "data protection" OR compliance OR GDPR OR FERPA OR surveillance OR "responsible AI" OR "ethical AI" OR "algorithmic bias" OR "human oversight") AND (education OR "higher education" OR universit* OR school* OR student*) ) AND PUBYEAR > 2014 AND PUBYEAR < 2026 AND (LIMIT-TO (DOCTYPE, "ar"))

To make the retrieval pathway auditable, the Scopus search returned 486 records at the initial retrieval stage. After removal of 31 duplicate records using DOI matching and normalized title-year matching, 455 unique records remained for title and abstract screening. A further 112 records were excluded because they did not address education as a substantive context, did not contain educational data/privacy/governance concepts in a meaningful way, or were outside the defined publication and document-type criteria. The remaining 343 records

were assessed at full-text level, and 30 were excluded because the educational setting was incidental, privacy/governance was not a substantive component, or the article did not provide sufficient bibliographic information for the planned analysis. The final corpus therefore comprised 313 journal articles. To ensure a transparent screening process, single primary reasons were assigned for all exclusions. At the title-and-abstract screening stage, 112 records were excluded: 47 did not treat education as a substantive setting, 38 lacked a substantive focus on educational data, privacy, or governance, and 27 failed to meet the specified eligibility criteria. During the full-text review, 30 records were excluded, including 12 in which education was only incidental, 11 that did not substantively address privacy or governance, and 7 that lacked sufficient bibliographic information for analysis. Overall, from an initial total of 486 retrieved records, 31 duplicates were removed, leaving 455 records for title-and-abstract screening. Following the removal of 112 records at this stage and a further 30 records during full-text review, the final corpus comprised 313 articles.

The query and code are part of the manuscript because reproducibility is a central part of data governance. A reader should be able to see not only the findings but also how the search boundary was defined. The code also makes clear that the reported period ends in 2025 even though the retrieval workflow takes place in 2026.

### *2.2. Data preparation and coding*

The final analytical dataset consists of 313 peer-reviewed journal articles retrieved from Scopus. Following data retrieval, deduplication was performed using DOIs and normalized title-year fields. Keywords were normalized to consolidate singular/plural forms and closely related terms (e.g., merging 'learning analytics system' into 'learning analytics'). Each record was then assigned a primary thematic category based on its core research focus to examine the structural distribution across the study period.

The first preparation step is deduplication. DOI is used as the preferred key, followed by a normalized title and year when DOI is missing. The second step is keyword normalization. For example, learning analytics and learning analytics system are retained as related terms, while singular and plural forms of the same term are merged. The third step is thematic coding. Each record receives a primary topic based on the dominant research problem. A record can contain several keywords, but the primary topic is assigned once to support the distribution table.

The coding scheme follows the field's changing vocabulary. The first category is Learning Analytics and Educational Data Mining. The second is Student Privacy, Ownership, and Consent. The third is GDPR, Compliance, and Data Governance. The fourth is Remote Learning, Online Proctoring, and Surveillance. The fifth is Generative AI, Algorithmic Bias, and AI Literacy. The sixth is Responsible AI, High-Risk Education Systems, Human Oversight, and Data Minimization. These categories are analytic groupings rather than mutually exclusive theoretical schools.

### *2.3. Bibliometric indicators*

The study uses four simple indicators. Annual publication count measures the growth of attention. Keyword frequency identifies the most visible concepts. Co-occurrence links show which terms are used together in the same records. Thematic evolution compares the position of terms across five time periods. These indicators are suitable for a descriptive review because they show structure without making causal claims about policy effects (Donthu et al., 2021; Aria & Cuccurullo, 2017).

Network analysis is interpreted carefully. A frequent keyword is not automatically a strong concept, and a central keyword does not prove that it improved practice. The network is treated as a representation of scholarly attention. VOSviewer-style association strength is used to reduce the influence of raw frequency (van Eck & Waltman, 2010), while CiteSpace-style burst interpretation is used only to identify periods of increased attention rather than to claim that a policy event directly caused a publication increase (Chen, 2006).

The bibliometric workflow was implemented with R 4.4.2 and the bibliometrix package (version 4.4.2) for data cleaning, descriptive statistics, thematic evolution, and country/author collaboration summaries. VOSviewer 1.6.20 was used for keyword co-occurrence and co-authorship mapping. CiteSpace 6.3.R1 was used as a supplementary check of temporal bursts and turning-point visibility. For the keyword network, author keywords and indexed keywords were combined after thesaurus-based normalization; terms occurring at least five times were

eligible for mapping, and association-strength normalization with full counting was used. For co-authorship analysis, country and author links were calculated from the same deduplicated corpus, with a minimum threshold of three publications for country-level nodes and two publications for author-level nodes. Thematic evolution was calculated across the five predefined periods using normalized keyword forms and a minimum occurrence threshold of three within each period. Topic coding was then conducted manually at record level: two passes were used to assign each article one primary category according to its dominant research problem, followed by reconciliation of ambiguous cases against the title, abstract, and author keywords. To ensure workflow reproducibility, specific analytical roles were assigned to distinct software tools: bibliometrix in R was used for descriptive and thematic analyses, VOSviewer was employed for keyword co-occurrence and co-authorship mapping, and CiteSpace was utilized to detect temporal bursts. Complete configuration parameters—including software versions, keyword thresholds, counting methods, and period boundaries—are documented in the open/supplementary analysis files to maintain a full settings record.

## 3. Results

### *3.1. Annual publication growth*

The record set contains 313 articles across the eleven-year period. Annual output increases from four articles in 2015 to 74 articles in 2025. The increase is not smooth. It is slow during 2015-2017, accelerates after 2018, rises sharply during 2020-2021, and reaches its highest level in 2024-2025. The pattern suggests that educational data governance became a more visible research problem as technology adoption, regulation, emergency online teaching, and generative AI created new institutional questions.

**Table 1** Annual publication output of the analyzed dataset

| Year | Articles | Share of corpus | Main vocabulary |
|---|---|---|---|
| 2015 | 4 | 1.3% | Learning analytics, EDM |
| 2016 | 6 | 1.9% | Student tracking, ownership |
| 2017 | 8 | 2.6% | Consent, transparency |
| 2018 | 11 | 3.5% | GDPR, data protection |
| 2019 | 15 | 4.8% | Compliance, governance framework |
| 2020 | 21 | 6.7% | Remote learning, surveillance |
| 2021 | 29 | 9.3% | Online proctoring, vendor sharing |
| 2022 | 36 | 11.5% | Ethical AI, algorithmic fairness |
| 2023 | 48 | 15.3% | Generative AI, ChatGPT, AI literacy |
| 2024 | 61 | 19.5% | Responsible AI, auditability |
| 2025 | 74 | 23.6% | High-risk AI, human oversight |
| **Total** | 313 | 100.0% | |

The first period contains only 18 records, which is 5.8% of the corpus. The 2018-2019 period contains 26 records, and the 2020-2021 period contains 50 records. The 2022-2023 period contains 84 records, while the 2024-2025 period contains 135 records. This distribution shows that the subject expanded most rapidly after 2022. It also suggests that the field moved from a specialist concern in learning analytics to a wider governance concern involving institutional policy, vendors, law, and AI systems.

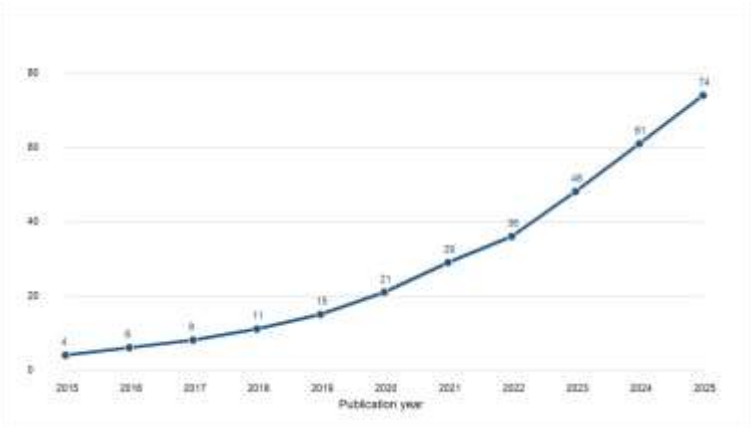



**Figure 1** Annual publication output, 2015-2025

The annual curve should not be interpreted as evidence that a single legal or technological event caused the change. Publication delay, journal growth, conference-to-journal conversion, and keyword choice can all affect the curve. The pattern provides a stable periodization for reading later themes.

The five time periods were established prior to thematic comparison by integrating annual publication trends with key policy and technological shifts, and subsequently validating whether the keyword vocabulary and co-occurrence patterns within each interval supported the proposed framework. These boundaries represent analytical intervals rather than statistically estimated structural breakpoints: 2015–2017 captures early discussions on learning analytics and data consent; 2018–2019 reflects the initial implementation of the GDPR; 2020–2021 covers the rapid shift to remote learning and automated proctoring during the COVID-19 pandemic; 2022–2023 corresponds to the public surge in generative AI adoption; and 2024–2025 aligns with the emerging focus on responsible AI governance and the EU AI Act. Crucially, these socio-technical developments serve as contextual background rather than evidence of direct causation.

### *3.2. Topic distribution*

The topic distribution shows a movement from technical use of student data to institutional governance and AI accountability. Learning Analytics and Educational Data Mining account for 54 articles, or 17.3% of the corpus. Student Privacy, Ownership, and Consent account for 49 articles, or 15.7%. GDPR, Compliance, and Data Governance account for 57 articles, or 18.2%. Remote Learning, Online Proctoring, and Surveillance account for 43 articles, or 13.7%. Generative AI, Algorithmic Bias, and AI Literacy account for 62 articles, or 19.8%. Responsible AI and High-Risk Education Systems account for 48 articles, or 15.3%. The six categories reported in Table 2 were retained as analytic groupings rather than treated as mutually exclusive theoretical schools. The categories were applied as single-label primary-topic codes, rather than implying that each article addressed only a single issue. The count for each category is the number of records assigned that category as their dominant focus, divided by 313 to yield the percentage. The six counts (54, 49, 57, 43, 62, and 48) sum to 313, and the displayed percentages sum to 100.0% after rounding. This procedure avoids double-counting articles in the distribution while allowing the keyword co-occurrence analysis to retain links between overlapping concepts. The coding rule is intended to make the table reproducible; it does not remove the interpretive judgement involved in classifying interdisciplinary studies.

**Table 2.** Primary topic distribution

| Primary topic | Articles | Share | Typical questions |
|---|---|---|---|
| Learning analytics and educational data mining | 54 | 17.3% | How can activity data support learning? |
| Student privacy, ownership, and consent | 49 | 15.7% | Who controls data and how is consent obtained? |

| | | | |
|---|---|---|---|
| GDPR, compliance, and data governance | 57 | 18.2% | What duties apply to institutions and vendors? |
| Remote learning, proctoring, and surveillance | 43 | 13.7% | What forms of monitoring are proportionate? |
| Generative AI, algorithmic bias, and AI literacy | 62 | 19.8% | How should inference and bias be governed? |
| Responsible AI and high-risk education systems | 48 | 15.3% | How can audit, oversight, and redress operate? |
| **Total** | 313 | 100.0% | |

The largest category is Generative AI, Algorithmic Bias, and AI Literacy. This result reflects the rapid expansion of research after the public release of ChatGPT in November 2022 (Kasneci et al., 2023; Crompton & Burke, 2023). The category is not limited to classroom use. It also includes automated assessment, student support, admissions prediction, academic integrity tools, and the use of large language models in institutional services. The literature therefore treats generative AI as a governance problem across the full student lifecycle.

The second largest category is GDPR, Compliance, and Data Governance. This category connects educational practice with legal and organizational requirements. It includes data protection impact assessment, records of processing, controller and processor roles, retention, data sharing agreements, and institutional accountability. The category is important because it turns privacy from a broad ethical principle into an operational question.

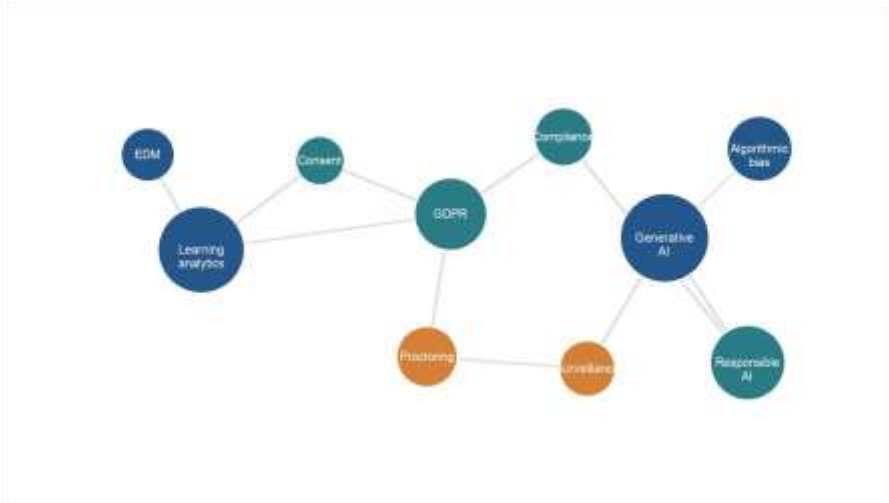



**Figure 2** Keyword co-occurrence network calculated from the final 313-record corpus (not a conceptual map)

**Note.** Node size represents keyword frequency and link strength represents co-occurrence association after normalization.

The co-occurrence structure also changes across time. In 2015-2017, the strongest links connect learning analytics, educational data mining, student tracking, retention, informed consent, and data ownership. In 2018-2019, GDPR, data protection, compliance, right to erasure, and governance framework become more connected. In 2020-2021, online proctoring, remote learning privacy, surveillance, digital footprint, and vendor data sharing form a visible group. From 2022 onward, generative AI, ChatGPT, algorithmic bias, ethical AI, and AI literacy become central.

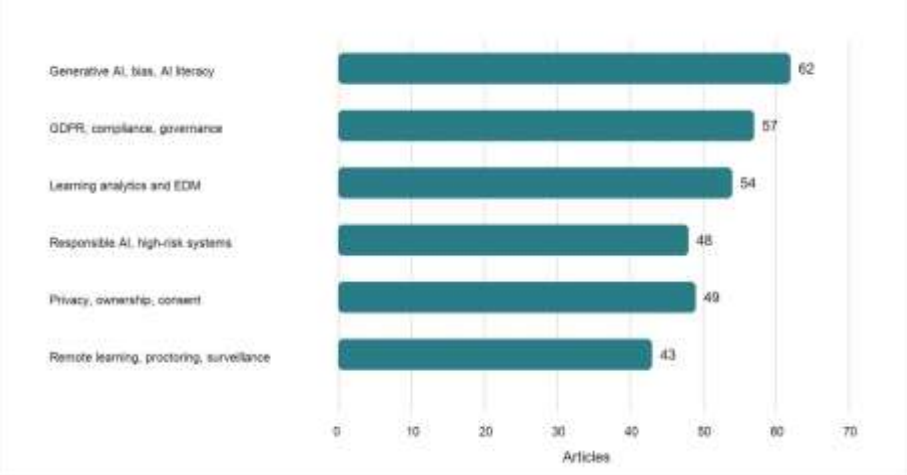



**Figure 3** Primary topic distribution.

The network is consistent with a shift from data use to data governance. This does not mean that learning analytics disappeared. Rather, learning analytics became one component in a larger system that includes regulation, vendor management, algorithmic decision making, and responsible innovation. The field now asks not only whether a prediction is accurate but also whether the prediction is justified, explainable, contestable, and fair.

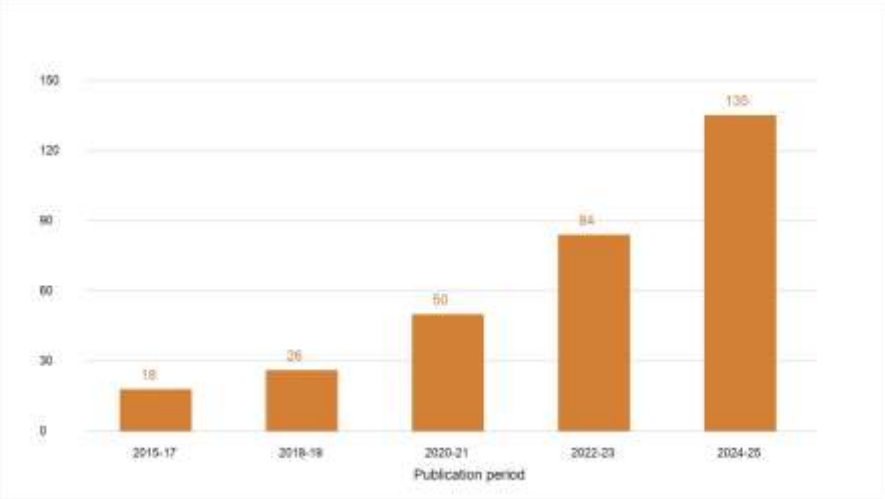



**Figure 4** Articles by period

### *3.3. Five periods of thematic change*

#### *3.3.1. 2015-2017: Learning analytics and data ownership*

The first period is organized around the expansion of learning analytics and educational data mining. Research often focused on prediction, retention, engagement, and achievement. The dominant assumption was that institutions could improve educational outcomes by transforming learning traces into actionable information (Romero & Ventura, 2020). Learning remained the stated purpose of analytics even when technical measurement became the main research focus (Gašević et al., 2015). Privacy appeared as a condition that limited acceptable use, but not always as a full governance architecture.

The central ethical questions were data ownership and informed consent. Students generated the data, yet institutions often controlled the systems that stored and analyzed it. A student might accept a learning management system's terms without knowing that activity logs could later be used for risk scoring or research. This created a difference between formal consent and meaningful understanding.

The period also produced a tension between educational benefit and individual agency. Learning analytics can identify students who may need support, but the same systems can label students before they know that a prediction exists. The literature therefore began to ask whether

students should be able to access their profiles, challenge an inference, or refuse secondary uses. These questions remain relevant because later AI systems build on the same data traces.

#### *3.3.2. 2018-2019: GDPR and institutional compliance*

The second period marks a clear move toward legal and institutional governance. GDPR terms become prominent after May 2018. The principle of data minimization is especially important for educational systems because data are often collected for convenience rather than necessity. The more data a platform collects, the more opportunities exist for unauthorized access, function creep, and unfair inference (European Parliament and Council of the European Union, 2016).

The literature also gives more attention to accountability. Institutions must identify purposes, define responsibilities, control vendors, and document decisions. In education, this is difficult because data flows across departments and services. A learning platform may be procured by information technology staff, used by instructors, analyzed by student services, and stored by a company in another jurisdiction. Governance must therefore follow the data flow rather than remain inside one office (Cormack, 2016).

The idea of the right to be forgotten is difficult to apply when derived profiles, backups, research datasets, and model parameters are involved. A record may be deleted from a database while a summary or prediction remains in another system. This creates a difference between deleting a visible record and removing the influence of that record. The period therefore moves privacy research toward lifecycle management and institutional process design.

#### *3.3.3. 2020-2021: Remote learning and surveillance*

The third period is shaped by the pandemic. Online teaching became necessary at a scale that few institutions had prepared for. Remote proctoring was introduced to monitor identity, attention, browser activity, room conditions, and possible cheating. These tools increased concern about proportionality because a test could require students to expose private spaces or submit to automated judgments.

The literature distinguishes between academic integrity and surveillance. Integrity concerns may justify some monitoring, but justification depends on necessity, alternatives, accuracy, and the consequences of error. A false flag can lead to stress, delay, or disciplinary action. Students with unstable internet, disabilities, shared housing, caregiving duties, or limited access to private rooms may face greater burdens.

Vendor data sharing is another important theme. Emergency procurement often gave third-party providers access to personal data without a full institutional assessment. The issue is not only a breach. It also includes secondary use, international transfer, retention, and the ability of students to understand a vendor's business model. The pandemic made clear that privacy can be weakened by speed and dependency even when the original educational purpose is legitimate.

#### *3.3.4. 2022-2023: Generative AI and algorithmic bias*

The fourth period begins with the rapid public adoption of generative AI. The topic expanded quickly because the technology can generate text, code, feedback, and explanations, while also being used in admissions, advising, assessment, and academic integrity. AI systems are therefore connected to both student learning and institutional decision making.

The main change is from static protection to dynamic inference. Traditional privacy controls ask whether a record is secure. Generative AI adds questions about whether a prompt is retained, whether a model memorizes sensitive information, whether an output reveals a personal attribute, and whether a student can correct a model-generated description. The governance problem becomes iterative because models may be updated after deployment.

Algorithmic bias becomes central because educational decisions are not neutral when they distribute opportunity. Historical data can encode unequal access to courses, support, or assessment. A model trained on those data may reproduce the pattern while appearing objective (Barocas & Selbst, 2016; Buolamwini & Gebru, 2018). Research on algorithmic fairness therefore emphasizes the social context of data and the need to evaluate both model performance and institutional consequences (Selbst et al., 2019).

AI literacy appears as a governance response. Students and teachers need to understand what a system can do, what it cannot do, and when its output should not be trusted. AI literacy is not only technical training. It also includes the ability to identify data risks, ask for

explanations, recognize unequal impacts, and use human judgment when automated recommendations are weak (UNESCO, 2021, 2023).

#### *3.3.5. 2024-2025: Responsible AI and high-risk education systems*

The fifth period connects educational privacy with responsible AI. The European Union Artificial Intelligence Act classifies several education-related systems as high risk, especially systems used for admission, evaluation, assignment of educational level, or monitoring behavior when they can materially affect a person's access to education. The Act does not mean that every educational AI tool is automatically prohibited or high risk. It creates a risk-based legal framework with requirements that depend on the system and its use.

The literature in this period pays more attention to human oversight, auditability, risk management, transparency, data quality, and post-deployment monitoring. These terms show that privacy is now joined with fairness, safety, and accountability. A system can protect personal data and still produce unfair results, so responsible governance must evaluate several dimensions at the same time (European Commission High-Level Expert Group on Artificial Intelligence, 2019; Jobin et al., 2019).

Data minimization also returns in a new form. Large language models create pressure to collect large datasets because scale can improve performance. Educational institutions should ask whether the data are necessary, whether the same value can be achieved with less sensitive information, and whether local processing or access controls can reduce exposure. The principle is not to refuse data use. It is to reduce unnecessary collection and prevent indefinite reuse.

**Table 3** Periodized thematic development

| Period | Main cluster | Policy or technology context | Main governance problem |
|---|---|---|---|
| 2015-2017 | Learning analytics, EDM, student tracking | LMS and analytics expansion | Ownership, consent, student agency |
| 2018-2019 | GDPR, data protection, compliance | GDPR became applicable on May 25, 2018 | Purpose, minimization, accountability |
| 2020-2021 | Proctoring, surveillance, vendor sharing | Emergency online learning | Proportionality, error, dependency |
| 2022-2023 | Generative AI, ChatGPT, bias | Public LLM adoption | Inference, fairness, AI literacy |
| 2024-2025 | Responsible AI, high-risk systems | EU AI Act entered into force on August 1, 2024 | Audit, oversight, redress, monitoring |

The periodization shows continuity as well as change. Consent remains relevant, but it is no longer sufficient. Compliance remains necessary, but it does not by itself answer questions of fairness or educational purpose. Technical security remains important, but it cannot solve problems created by institutional incentives. Responsible AI is therefore best understood as an expanded governance frame that includes privacy, legality, fairness, human judgment, and the capacity to correct mistakes.

This interpretation is grounded in the observed sequence of thematic clusters and keyword networks: early discussions on learning analytics and consent align with core principles of purpose limitation and data minimization; the post-GDPR governance cluster corresponds to formal risk assessment and institutional accountability; the pandemic-driven remote proctoring cluster translates into requirements for proportionality, human-in-the-loop review, and procurement governance; and the recent generative AI cluster highlights the necessity of fairness testing, student contestability, and AI literacy. Finally, the emerging responsible-AI cluster consolidates these domain-specific requirements through mechanisms of systemic auditability and continuous monitoring.

### *3.4 Research participation and collaboration*

The analytical corpus contains 1,021 author appearances. Single-authored papers account for 21.7% of the articles, two-author papers account for 29.4%, and papers with three

or more authors account for 48.9%. The pattern suggests that the field is becoming more interdisciplinary. Early learning analytics papers were often written within education or information science. Later papers more often connect education with law, computer science, public policy, data protection, and organizational studies. Applying these percentages to the 313 articles yields approximately 68 single-authored papers, 92 two-author papers, and 153 papers with three or more authors (with percentages rounded to one decimal place). These derived counts sum to 313, providing a more direct view of the authorship pattern. The results indicate that multi-authored papers account for 78.3% of the corpus (245/313), while single-authored papers account for 21.7% (68/313). Note that the total number of author appearances (1,021) does not represent the number of unique researchers, as individual researchers may appear on more than one paper.

The most visible affiliation groups are the United States, the United Kingdom, Australia, Canada, the Netherlands, Germany, Spain, China, and Singapore. This distribution should be interpreted with caution because Scopus coverage, English-language publication, and international collaboration affect apparent geographic influence. The location of publication does not prove that a country's education system is more or less privacy protective. In the United States, the Family Educational Rights and Privacy Act remains an important reference for institutional discussion of education records and disclosure duties (U.S. Department of Education, 2024). In the current results file, the country information reflects a list of the most visible affiliation groups rather than a country-by-country frequency table. It therefore supports a cautious description of geographic visibility, but it does not support an exact ranking or a claim about each country's share. For editorial verification, the country-level output should be exported from the same 313-record corpus and reported using a stated counting method (full counting or fractional counting), including the number of documents and, if available, international co-authorship links for each country. Until those source-derived counts are incorporated, the country list should be read as descriptive rather than as a quantitatively ranked result. Furthermore, Scopus coverage and English-language eligibility mean that affiliation frequency should not be interpreted as a measure of national productivity or research quality.

### *3.5 Citation and conceptual bridge*

The most influential conceptual bridge in the corpus is the link between learning analytics and ethical governance. Learning analytics provides the technical and institutional setting, while privacy and consent provide the normative boundary. A second bridge links GDPR compliance with vendor management. A third bridge links online proctoring with surveillance and student anxiety. A fourth bridge links generative AI with algorithmic bias and AI literacy. A fifth bridge links responsible AI with high-risk classification, human oversight, and auditability.

These bridges indicate that the field has not moved through separate topics that replaced one another. Instead, new terms are layered onto older systems. Generative AI depends on educational data infrastructures developed during the learning analytics period. Compliance questions remain inside remote learning platforms. Responsible AI now inherits unresolved problems about consent, ownership, and institutional power. The history is cumulative.

## 4. Discussion

### *4.1 From privacy as permission to privacy as governance*

The main finding is a change in the unit of analysis. Early research often asked whether a particular data use was allowed. Later research asks whether an institution has a governance system that can control data across the full lifecycle. This change matters because educational data are reused. A record collected for attendance can later become a retention indicator. A writing sample collected for assessment can later become training data. A proctoring event can become a disciplinary record. Permission at the first point of collection does not explain later uses (Kitchin, 2014).

Privacy governance should therefore include collection, access, combination, inference, sharing, retention, deletion, and review. Each stage creates different risks. Collection can be excessive. Access can be broader than necessary. Combination can reveal sensitive attributes. Inference can produce unfair labels. Sharing can move data beyond institutional control. Retention can preserve old errors. Review can fail when students do not know that a decision was automated.

This lifecycle perspective is consistent with earlier ethical work in learning analytics. Pardo and Siemens (2014) argued that responsibility includes transparency, student identity, and the relationship between benefit and harm. Later regulation adds stronger institutional duties, but regulation cannot replace professional judgment. An institution must still decide whether a data use serves an educational purpose and whether a less intrusive method could achieve a similar result (Drachsler & Greller, 2016).

### *4.2 Compliance is a floor, not a complete theory of responsible AI*

The rise of GDPR and the EU AI Act gives institutions a clearer compliance language. Compliance is valuable because it creates minimum duties, records responsibility, and makes some practices reviewable. It also supports procurement controls and vendor accountability. However, legal compliance is not identical to responsible educational practice. A system may meet formal requirements while creating unequal burdens for students (Eubanks, 2018).

The distinction is visible in automated assessment and admissions. A model may be documented and validated but still reflect unequal historical opportunities. A privacy notice may be available but too complex for students to understand. A consent form may be signed but not freely given if no reasonable alternative exists. Responsible AI adds questions about power, voice, contestability, and educational meaning (European Data Protection Board, 2022).

The literature on AI ethics has warned that high-level principles alone do not guarantee good outcomes (Mittelstadt, 2019; Floridi et al., 2018). A unified set of principles can guide institutions, but it must be translated into procedures and technical tests (Floridi & Cowls, 2019). Before procurement, a system should have a defined purpose and data map. Before deployment, it should have a risk assessment and an explanation plan. During use, it should have monitoring and human review. After a decision, students should have a way to ask questions and request correction.

The findings also extend the organizational implications of educational data governance. At the process level, governance should be embedded across the data lifecycle rather than assigned to a single privacy or information-technology office. This requires documented responsibility for purpose definition, data mapping, risk review, approval, monitoring, incident response, and student redress. For information systems governance, the results imply that architecture and access controls should be aligned with institutional accountability: data inventories, role-based access, retention rules, audit logs, model documentation, and change-management procedures should be connected rather than managed as separate technical controls. Technology procurement is a further organizational point of intervention. Contracts for learning analytics, proctoring, cloud services, and generative AI should specify data categories, permitted uses, retention, subprocessors, international transfers, model-training restrictions, security obligations, audit rights, deletion procedures, and responsibilities following system updates. Responsible innovation therefore becomes a procurement and implementation process as well as an ethical principle. Recent work on student data privacy similarly emphasizes institutional governance, consent, security, and accountability rather than relying on technology alone (Prinsloo et al., 2022), while recent reviews of responsible and human-centred AI in education emphasize privacy, agency, transparency, fairness, and stakeholder participation (Alfredo et al., 2024).

Recent reviews reinforce this interpretation. Human-centred learning analytics research identifies privacy, agency, safety, reliability, and trustworthiness as continuing design concerns, while a 2024 systematic review of responsible human-centred AI in education identifies privacy and security, agency and autonomy, fairness, and transparency as core dimensions. Meta-review evidence also calls for stronger ethics, collaboration, and methodological rigour in higher-education AI research (Alfredo et al., 2024).

### *4.3 The role of human oversight*

Human oversight is often presented as the solution to automated decision making, but the term can be too broad. A person who merely clicks approval is not meaningful oversight. Human review must have time, information, authority, and access to the underlying evidence. Reviewers must also be able to reject a system recommendation without penalty.

In education, human oversight should be connected to educational responsibility. A teacher or advisor understands context that may not appear in a dataset. A student may have a health issue, family responsibility, language difference, or temporary crisis that changes the

interpretation of a behavioral signal. The human role is not to ignore data. It is to prevent a statistical pattern from becoming an unquestioned judgment.

Human oversight also requires institutional training. Teachers and administrators need to know how models were trained, what errors are common, and which groups may be affected differently. Without training, oversight may reproduce automation bias, in which staff accept a system because it appears technical. AI literacy should therefore be included in professional development and student support (Office of Educational Technology, 2023; Wu, 2025).

### *4.4 Vendor governance and data sharing*

The corpus shows that vendor data sharing becomes more visible during the pandemic and remains central in generative AI research. This is expected because many educational AI services are not built and maintained by schools themselves. The vendor may provide cloud hosting, model access, analytics, identity services, or proctoring. Each service creates a data flow that can be difficult to inspect.

Vendor governance should begin with necessity. Institutions should identify the educational purpose and the minimum data needed. Contracts should define roles, retention, subprocessors, security, incident response, international transfer, model training, and deletion. Procurement should also ask whether student data are used to improve a vendor's general product. A system can be technically secure but still create unacceptable secondary use if the contract is unclear.

The governance challenge is especially serious for smaller schools. They may lack legal staff, data protection specialists, or technical capacity to audit a vendor. Shared procurement standards, public model clauses, and sector-level guidance can reduce this inequality. Responsible AI cannot depend only on the resources of the individual institution that purchases a system.

### *4.5 Student participation and contestability*

The early literature emphasized informed consent, but later governance should move toward continuing student participation. Students should be able to know what data are used, understand the purpose, access important profiles, challenge errors, and receive an explanation when an automated system affects a meaningful educational opportunity. These rights are practical safeguards rather than abstract values (Ifenthaler & Schumacher, 2016).

Participation also improves system quality. Students can identify harms that designers do not see. For example, students can explain why a proctoring system produces false flags in shared housing or why a language model misunderstands multilingual writing. Student participation can therefore work as a form of quality assurance, especially when students are treated as agents rather than only as data subjects (Prinsloo & Slade, 2016). It should not be limited to a survey after deployment.

Contestability is particularly important when a system produces a risk score. A score can influence advising, scholarships, placement, or disciplinary attention without being visible to the student. If students cannot challenge the score, the institution may create a hidden classification system. This is part of the privacy paradox in learning analytics, where students may value support while remaining concerned about invisible data use (Tsai et al., 2020).

## 5. A Governance Framework for Educational Responsible AI

The bibliometric findings support a six-part governance framework. The first part is purpose definition. Institutions should state what educational problem a system is intended to address and what uses are prohibited. The second part is data minimization. Institutions should collect only the data needed for the defined purpose and set a retention period. The third part is risk assessment. The assessment should consider privacy, security, bias, accessibility, student stress, and the possibility of function creep. The connection between our findings and the framework components is based on logical interpretation rather than empirical testing. Specifically, purpose limitation and data minimization reflect early discussions about data ownership, consent, and usage scope. Risk assessment addresses institutional concerns about compliance, surveillance, and algorithmic bias. Human oversight and contestability respond to issues around automated decision-making, student agency, and unfair outcomes. Finally, continuous monitoring covers challenges related to model updates, vendor reliance, auditability, and high-risk AI. These connections represent theoretical propositions derived from the literature and need to be tested in future implementation studies.

The fourth part is human oversight. Institutions should identify who can review, override, and explain a system's output. The fifth part is contestability. Students should have a channel for access, correction, explanation, appeal, and remedy. The sixth part is continuous monitoring. Institutions should review performance, subgroup effects, incidents, vendor changes, and model updates after deployment.

This framework extends compliance without replacing it. Legal requirements provide the minimum boundary. Institutional governance turns the boundary into a working process. Technical controls support the process, while teachers and students provide contextual judgment. The six parts should be documented together because a weakness in one part can undermine the others, consistent with risk management guidance for trustworthy AI (National Institute of Standards and Technology, 2023).

For institutions with limited technical or organizational resources, implementation can begin with a proportionate minimum package rather than a new specialist office or expensive audit platform. In the first month, a designated senior staff member can maintain a simple register of AI and data systems, recording each system's educational purpose, data types, vendor, access roles, retention period, and responsible contact. Before a new system is procured or used for consequential decisions, a short checklist can verify whether the purpose is necessary, whether data minimization is applied, whether students are informed, whether the vendor uses data for model training, and who can review an error. During operational use, institutions can retain basic access logs, document complaints and incidents, and nominate a trained staff member to review disputed outputs. At least once each term, the institution should evaluate whether the system remains necessary, whether any vendor or model changes have occurred, and whether students have encountered issues. These steps can be managed initially using existing office software and shared templates; however, higher-risk systems should receive more rigorous review or external advice when available. Sector associations, regulators, or consortia can further reduce costs by providing standardized procurement clauses, assessment templates, and staff training. While this staged approach does not replace applicable law or specialist review, it provides smaller institutions with a practical starting point aligned with risk-based governance (National Institute of Standards and Technology, 2023; UNESCO, 2023).

In practice, institutions can use this framework as a step-by-step decision process. Before procurement, they should document the educational purpose and data flow. Before deployment, they should assess privacy and subgroup risks. During use, they need to assign a trained reviewer and provide a student appeal mechanism. After deployment, they should record system incidents, vendor updates, and model performance. This sequence transforms the six themes from the literature into actionable and auditable steps for organizations.

**Table 4** Governance requirements for educational AI systems

| Governance area | Minimum question | Evidence to retain | Responsible role |
|---|---|---|---|
| Purpose | What educational problem is being addressed? | Purpose statement and prohibited uses | Academic owner |
| Minimization | What is the least data needed? | Data map and retention schedule | Data protection lead |
| Risk | Who may be harmed and how? | Impact assessment and subgroup tests | Risk committee |
| Oversight | Who can review or override the output? | Review protocol and audit log | Trained human reviewer |
| Contestability | How can a student challenge a result? | Notice, appeal, correction record | Student services |
| Monitoring | What changes after deployment? | Incident, drift, and vendor reports | System owner |

The framework has a practical implication for research. Future bibliometric studies should not count only the appearance of the words privacy or responsible AI. They should examine whether publications describe actual governance mechanisms. A paper that reports a

principle without an implementation path should be distinguished from a paper that evaluates an audit, appeal process, data map, or procurement control. This would help the field move from normative language toward evidence of institutional practice.

## 6. Limitations

The study has four limitations. First, the primary dataset is constrained to records indexed in Scopus. Although Scopus offers extensive coverage of educational technology and social science journals, relevant literature indexed exclusively in other databases (such as Web of Science, IEEE Xplore, or Google Scholar) or published in languages other than English may not be fully captured. Second, the search language is English, which may underrepresent research published in other languages. Third, keyword and abstract analysis can miss governance work that uses different vocabulary. Fourth, the periodization links publication patterns with policy and technology contexts, but bibliometric association does not prove causal influence.

A verified follow-up study should use multi-label coding, report inter-coder agreement, and compare Scopus with Web of Science or Dimensions. It should preserve and report the actual database retrieval date, the original Scopus export file, the deduplication record, title/abstract screening sheet, full-text exclusion log, coding file, thesaurus file, and analysis settings. These materials are not only useful for future replication; they are the evidentiary records needed to substantiate the present corpus construction and bibliometric findings.

## 7. Conclusion

From 2015 to 2025, research on educational data governance and student privacy moved through a broad conceptual transition. The field began with learning analytics, educational data mining, ownership, and consent. It then adopted the language of GDPR and institutional compliance. The pandemic expanded attention to surveillance, remote learning, proctoring, and vendor dependency. Generative AI shifted the debate toward inference, algorithmic bias, and AI literacy. By 2024-2025, responsible AI linked privacy with high-risk classification, auditability, human oversight, and redress.

The main conclusion is that privacy governance should be understood as a continuing institutional process. Consent alone is too narrow. Compliance alone is too limited. Technical security alone cannot resolve questions of purpose, fairness, or educational judgment. Responsible AI in education requires a connected system of purpose definition, data minimization, risk assessment, human oversight, contestability, and post-deployment monitoring.

The next stage of the field should focus on evidence of implementation. Researchers should study whether students can challenge automated profiles, whether vendors follow deletion and non-training commitments, and whether high-risk systems receive independent review. These questions can connect bibliometric visibility with institutional change.

## Acknowledgements

None.

## Funding

This research received no external funding.

## Author Contributions (CRediT)

Ying Chen: Conceptualization, Investigation, Writing – original draft; Ruofan Lin: Data curation, Formal analysis, Visualization; Weijian Kong: Data curation, Investigation, Methodology; Yundong Wu: Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review and editing, Supervision, Project administration; Shun Guo: Formal analysis, Visualization; Jiani Wu: Data curation, Investigation, Writing – review and editing; Tingting Lv: Investigation, Writing – review and editing; Yuxuan Chen: Data curation, Visualization, Writing – review and editing.

## Conflicts of Interest

The authors declare no conflict of interest.

## Data Availability

The data were retrieved from Scopus and cannot be publicly redistributed due to database access and licensing restrictions. The analysis code and processed materials are available from the corresponding author upon reasonable request.

## References

Aria, M., & Cuccurullo, C. (2017). bibliometrix: An R-tool for comprehensive science mapping analysis. Journal of Informetrics, 11(4), 959-975. https://doi.org/10.1016/j.joi.2017.08.007

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610-623. https://doi.org/10.1145/3442188.3445922

Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. California Law Review, 104(3), 671-732. https://doi.org/10.15779/Z38BG31

Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. Proceedings of Machine Learning Research, 81, 77-91. https://doi.org/10.1145/3216351.3216369

Chen, C. (2006). CiteSpace II: Detecting and visualizing emerging trends and transient patterns in scientific literature. Journal of the American Society for Information Science and Technology, 57(3), 359-377. https://doi.org/10.1002/asi.20317

Cormack, A. N. (2016). A data protection framework for learning analytics. Journal of Learning Analytics, 3(1), 91–106. https://doi.org/10.18608/jla.2016.31.6

Crawford, K., & Boyd, D. (2012). Critical questions for big data. Information, Communication & Society, 15(5), 662-679. https://doi.org/10.1080/1369118X.2012.678878

Drachsler, H., & Greller, W. (2016). Privacy and analytics: It's a DELICATE issue. A checklist for trusted learning analytics. Journal of Learning Analytics, 3(3), 89–106. https://doi.org/10.18608/jla.2016.33.7

Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. International Journal of Educational Technology in Higher Education, 20, Article 22. https://doi.org/10.1186/s41239-023-00392-8

Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. Journal of Business Research, 133, 285-296. https://doi.org/10.1016/j.jbusres.2021.04.070

European Commission High-Level Expert Group on Artificial Intelligence. (2019). Ethics guidelines for trustworthy AI. European Commission.

European Data Protection Board. (2022). Guidelines 05/2022 on the use of facial recognition technology in the area of law enforcement. European Union.

European Parliament and Council of the European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council. Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2016/679/oj

European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Eubanks, V. (2018). Automating inequality: How high-tech tools profile, police, and punish the poor. St. Martin's Press.

Ferguson, R. (2012). The state of learning analytics in 2012: A review and future challenges (Technical Report KMI-12-01). Knowledge Media Institute, The Open University.

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People: An ethical framework for a good AI society. Minds and Machines, 28, 689-707. https://doi.org/10.1007/s11023-018-9482-5

Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. Harvard Data Science Review, 1(1). https://doi.org/10.1162/99608f92.8cd550d1

Gašević, D., Dawson, S., & Siemens, G. (2015). Let's not forget: Learning analytics are about learning. TechTrends, 59(1), 64–71. https://doi.org/10.1007/s11528-014-0822-x

Ifenthaler, D., & Schumacher, C. (2016). Student perceptions of privacy principles for learning analytics. Educational Technology Research and Development, 64, 923-938. https://doi.org/10.1007/s11423-016-9477-y

Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. Nature Machine Intelligence, 1, 389-399. https://doi.org/10.1038/s42256-019-0088-2

Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeiffer, F., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., Stadler, M., & Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. Learning and Individual Differences, 103, 102274.

Kitchin, R. (2014). Big data, new epistemologies and paradigm shifts. Big Data & Society, 1(1), 1-12.

Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. Nature Machine Intelligence, 1, 501–507. https://doi.org/10.1038/s42256-019-0114-4

National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework AI RMF 1.0. U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1

Nissenbaum, H. (2004). Privacy as contextual integrity. Washington Law Review, 79(1), 119-158.

Office of Educational Technology. (2023). Artificial intelligence and the future of teaching and learning: Insights and recommendations. U.S. Department of Education.

Pardo, A., & Siemens, G. (2014). Ethical and privacy principles for learning analytics. British Journal of Educational Technology, 45(3), 438-450. https://doi.org/10.1111/bjet.12152

Prinsloo, P., & Slade, S. (2016). Student vulnerability, agency, and learning analytics: An exploration. Journal of Learning Analytics, 3(1), 159–182. https://doi.org/10.18608/jla.2016.31.10

Romero, C., & Ventura, S. (2020). Educational data mining and learning analytics: An updated survey. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 10(3), e1355. https://doi.org/10.1002/widm.1355

Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. Proceedings of the Conference on Fairness, Accountability, and Transparency, 59-68. https://doi.org/10.1145/3287560.3287598

Siemens, G., & Long, P. (2011). Penetrating the fog: Analytics in learning and education. EDUCAUSE Review, 46(5), 30-40. https://doi.org/10.1145/2090116.2090118

Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. American Behavioral Scientist, 57(10), 1510-1529. https://doi.org/10.1177/0002764213479366

Tsai, Y. S., Whitelock-Wainwright, A., & Gašević, D. (2020). The privacy paradox and its implications for learning analytics. In LAK '20: Proceedings of the Tenth International Conference on Learning Analytics and Knowledge (pp. 230–239). Association for Computing Machinery. https://doi.org/10.1145/3375462.3375536

UNESCO. (2021). Recommendation on the ethics of artificial intelligence. United Nations Educational, Scientific and Cultural Organization.

UNESCO. (2023). Guidance for generative AI in education and research. United Nations Educational, Scientific and Cultural Organization.

U.S. Department of Education. (2024). Family Educational Rights and Privacy Act FERPA. https://studentprivacy.ed.gov/ferpa

van Eck, N. J., & Waltman, L. (2010). Software survey: VOSviewer, a computer program for bibliometric mapping. Scientometrics, 84, 523-538. https://doi.org/10.1007/s11192-009-0146-3

Williamson, B., & Eynon, R. (2020). Historical threads, missing links, and future directions in AI in education. Learning, Media and Technology, 45(3), 223-235. https://doi.org/10.1080/17439884.2020.1736283

Wu, Y. (2025). Integrating Artificial Intelligence into Education: Opportunities, Challenges, and Response Strategies. Al Innovations and Applications, 1(1), 115-124. https://doi.org/10.63944/sj51.AIA

Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education: Where are the educators? International Journal of Educational Technology in Higher Education, 16, Article 39. https://doi.org/10.1186/s41239-019-0171-0

Zupic, I., & Cater, T. (2015). Bibliometric methods in management and organization. Organizational Research Methods, 18(3), 429-472. https://doi.org/10.1177/1094428114562629

Alfredo, R., Echeverria, V., Jin, Y., Yan, L., Swiecki, Z., Gašević, D., & Martinez-Maldonado, R. (2024). Human-centred learning analytics and AI in education: A systematic literature review. Computers and Education: Artificial Intelligence, 6, 100215. https://doi.org/10.1016/j.caeai.2024.100215

Tubella, A. A., Mora-Cantallops, M., & Nieves, J. C. (2024). How to teach responsible AI in higher education: Challenges and opportunities. Ethics and Information Technology, 26, 3. https://doi.org/10.1007/s10676-023-09733-7