

THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

FROM INNOVATION TO REGULATION: ETHICAL GOVERNANCE OF AI IN TEACHING AND LEARNING

Sania Rehan (Principal Author)

PhD Scholar, Lahore College for Woman University, Lahore.

Kainatreh21@gmail.com

Zoya Faisal (Corresponding Author)

Lecturer, FAST School of Management, National University of Computer and Emerging
Sciences (FAST-NUCES). zoya.faisal@nu.edu.pk

Dr. Farzana Yousaf (Co-Author)

Assistant Professor of Education, University of Education, Bank Road Campus

Lahore. farzana.yousaf@ue.edu.pk

Fatima Faisal (Co-Author)

Independent Researcher, Institute of Business Administration (IBA), Computer Science at the
School of Mathematics and Computer Science (SMCS), Karachi,

Pakistan. fatimafaisal223@gmail.com


Abstract

This study examines the policy and ethical issues surrounding the implementation of Artificial Intelligence (AI) in higher education teaching and learning. With AI tools quickly finding their way into the classroom, which has been facilitated by AI in examining assessment, tutoring, feedback, and instructional design, they have also brought about the privacy, bias, transparency, academic integrity, and the changing role of educators. The aim of the study is to chart these ethical emerging issues and determine the strengths and weaknesses in the existing institutional and international policies regulating the utilization of AI in teaching. Peer-reviewed articles, institutional guidelines, and global policy documents published since 2020 to 2025 were used to carry out a qualitative scoping review. The PRISMA criteria were used to filter sources. Trends were classified by thematic analysis to categorize them into five areas, data governance, fairness and bias, pedagogical autonomy, student well-being, and accountability in AI-assisted decision-making. University and accreditation policy documents, as well as policy documents of various organizations like UNESCO, OECD, and EDUCAUSE, were compared in order to evaluate their coherence and readiness to implement. The findings suggest that although educational institutions of higher learning are becoming more responsive to the risks of ethical issues, policies are usually general, inconsistent, and responsive. The transparency of algorithms and their responsible use in grading and assessment is not well-advised, and there are no systems that can be used to track the long-term educational effect of AI. The faculty and students are not always clear on what they can do with the use, and this has resulted in uneven adoption. Better systems are required to provide responsible uses of AI that safeguard academic values and encourage equity and meaningful human supervision. The research identifies the necessity of standardized rules, policy creation, and constant training of teachers to make recommendation of ethical and sustainable integration of AI.

Keywords: Accountability, Bias, Ethics, Governance, Policy, Transparency, Trust

Introduction

Artificial intelligence (AI) has rapidly transformed the higher education sector, reshaping how educators design curricula, assess student learning, and support academic decision-making (Batoool, et al., 2024). To enhance the efficiency of learning and the accessibility of personalised education, universities are increasingly implementing AI-powered solutions, including automated grading systems and intelligent tutoring systems, predictive analytics dashboards and generative language models (Holmes et al., 2022). With the implementation of AI technologies at the level of higher education systems globally, the pedagogical market has become less instructor oriented, and more data oriented and processed by algorithms through the methods of teaching and learning (Luckin et al., 2016). Although these tools have potentials to streamline the assessment process, improve feedback loops, and help students to achieve



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

success, they also give rise to complicated ethical dilemmas stemming from the issues of privacy, fairness, transparency, accountability, and the changing role of the educator (Williamson and Eynon, 2020).

These large-scale generative models like ChatGPT, Gemini, and Claude further subject higher education to the impact of AI adoption by offering unprecedented content generation, writing assistance, and formative feedback functionality (Kasneji et al., 2023). Such systems have posed some important concerns regarding the integrity of academic system, reliability of model results, and the development of cognition skills in students (Makhdum & Mian, 2012). Although universities look forward to the opportunities of AI in improving the efficiency of instruction, the issues of the algorithmic bias, control of information, and the possibility of deskilling of academic labor remain influential in shaping the policy (Zawacki-Richter et al., 2019). What models like the Guidance to AI in Education by UNESCO emphasize is that ethical use of AI, being transparent and equitable, must be strongly governed and clearly defined by the institutional policies (UNESCO, 2023). However, most institutions of higher education do not have detailed, practical guidelines that react to the pace and the dynamism of AI implementations (Faisal, et al., 2024).

Here, it is important to know the ethical basis and policy frameworks on AI in higher education. The pedagogical and ethical implications are increasingly becoming higher as institutions incorporate AI in the curriculum design process, learning analytics, assessment tool, and student support services. By mapping such developments, the educators, administrators, and policymakers are able to spot the gaps in the current practices and imagine responsible ways forward. This paper analyses the legal and institutional principles that influence the application of AI in teaching and learning in higher education in terms of challenges in governance, risks, and new regulatory practices (Faisal, et al., 2024).


Significance

This study is important because there is a pressing issue that institutions of higher learning need to formulate coherent policies that are grounded on ethics to guide the use of AI. The AI technologies are being integrated into the fundamental academic procedures, such as plagiarism detection, automated grading, predicting student behavior and creating individual study plans (Pardo and Siemens, 2019). Nevertheless, lot of the policies are either fragmented or too general and do not give enough indication on accountability, transparency and data protection. Bias in predictive systems, opaque algorithmic decision-making, and misuse of student data are examples of ethical risks that may impact marginalized learners more than others and decrease the trust people have in institutions (Kitto and Villachica, 2023).

In addition, faculty do not know when and how AI tools can be applied in a responsible way in the teaching process, and students are worried about privacy, fairness, and exploitation of their information by automated systems (Ifenthaler and Schumacher, 2019). With the advent of the fast-moving generative AI, universities are under pressure to implement technologies without even comprehending their consequences (Makhdum & Mian, 2012). Devoid of sound policy frameworks, the institutions of higher education will end up implementing AI in a manner that undermines pedagogical integrity, ethical principles, and the autonomy of the learner.

The study is relevant since it gives a systematic discussion of the ethical issues and policy responses in global higher-education settings. It can be used to create more transparent, fairer, and more responsible models of teaching with the help of AI by streamlining current research, institutional policy, and global standards.

Problem Statement



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

Although the use of AI tools is becoming part of teaching in higher educational institutions, the ethical implications and the problem of governance have not been adequately discussed. Assessment, content generation, learning analytics, and automated feedback policies are not clearly defined in many institutions, and best practices are not agreed upon by educators and learners. The available frameworks are too general, lack operative particularity, lack any response to both algorithmic bias and transparency, and have no sufficient protection of student data or academic integrity. This disconnects between quick technological adoption and slow policy formulation poses ethical weaknesses and uneven application of the same in institutions (Makhdum & Mian, 2012).

Objective

- To analyze the ethical considerations associated with AI use in higher-education teaching, including issues of privacy, fairness, transparency, and accountability.
- To examine existing institutional and international policy frameworks guiding AI use in higher-education teaching.
- To identify gaps, inconsistencies, and areas requiring further policy development to support responsible AI integration.

Research Questions

- What are the major ethical challenges associated with using AI for teaching and learning in higher education?
- How do current institutional and international policy frameworks address the ethical risks of AI-enabled teaching?
- What gaps exist in current policies, and how can they be strengthened to support responsible AI adoption?

Gaps


There are some crucial gaps that exist in existing literature. To begin with, many research works investigate the AI use in education; however, a smaller number of them critically evaluate ethical considerations of AI application in teaching and assessment (Zawacki-Richter et al., 2019). Second, AI policies vary among institutions, many of them do not explicitly specify how to be transparent, how to govern data, or what instructional uses may be acceptable (UNESCO, 2023). Third, the vast majority of institutional structures are reactive and have been created in reaction to the rapid adoption of AI instead of through active ethical reflection (Kitto & Villachica, 2023). Fourth, research tends to be technical effectiveness-oriented instead of long-term pedagogical and equity consequences of such a practice (Holmes et al., 2022). Lastly, is lack of synthesis of intersection of global guidelines, institutional practices and ethical principles to guide policy formulation. This research fills such gaps by bringing together the analysis of ethics and policy review in the context of higher education.

Literature Review

AI Adoption in Higher Education

The introduction of AI in higher education has grown quickly, and such technologies as automated assessment systems, adaptive learning platforms, learning analytics dashboards, natural language generation tools, have become widespread. According to Holmes et al. (2022), AI technologies are becoming more common to tailor the learning process and offer immediate feedback. On the same note, Luckin et al. (2016) claim that AI helps to change the educational ecosystem by enhancing the abilities of instructors and data-driven decision-making.

The systematic review of Zawacki-Richter et al. (2019) indicates that the majority of AI-implemented solutions in higher education aim at administrative efficiency, data-driven



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

analytics, and student engagement prediction, which leads to the issue of algorithmic fairness and transparency. In the meantime, the development of generative AI has transformed the academic discourse, allowing students to produce essays, resolve complex problems, and synthesize information, which has made universities rethink the frameworks of academic integrity (Kasneci et al., 2023).

Ethical Considerations in AI-Enabled Teaching

Numerical ethical issues with the use of AI in higher education are usually categorized into four areas:

privacy, bias, transparency, and accountability.

The privacy issues are brought forth by a large amount of data collection such as behavioral records, demographic data, and test patterns that are used to train AI models. Williamson and Eynon (2020) note that surveillance-resembling learning analytics can change the behavior of students and destroy trust. In the study conducted by Ifenthaler and Schumacher (2019), students report that they feel uncomfortable with the idea of monitoring of their personalized learning systems.

Another major concern is prejudice. Training either on historical or unrepresentative data may support inequities in evaluation or academic predictions by using algorithmic systems. Pardo and Siemens (2019) caution that predictive modelling can have a systematic detrimental effect on disadvantaged students.

The transparency issues revolve around the black box of most AI algorithms. The faculty tends to be unaware of how automated grading or recommendation systems are made (Holmes et al., 2022).

The question of accountability is connected with how to assign responsibility in case AI tools perform missteps or produce possibly harmful results. UNESCO (2023) emphasizes the fact that AI-assisted teaching should have human control as a key element to maintain the agency of academic organizations.


Policy Frameworks for AI in Higher Education

Although the interest of AI governance in institutions is growing, the policies are still unequal. The UNESCO (2023) guidance on AI in Education offers the most commonly used international framework and human-focused AI, transparency, fairness, and effective data governance. Guidelines provided by OECD (2021) also recommend risk-based strategies and institutional responsibility measures.

But the policies in universities are very different. There are accounts of institutions explaining procedures of generative AI use in coursework, assessment, and faculty duties, whereas there are also general principles with no specific operational instructions (Kitto and Villachica, 2023). According to Zawacki-Richter et al. (2019), the majority of the policies focus on innovation rather than on governance, and the ethical considerations are not well-developed. Moreover, numerous frameworks do not provide any mechanisms of overseeing the long-term pedagogical effects of AI tools or guaranteeing equity within the social-economic/cultural groupings. Such lack of specificity tends to make instructors be unable to understand how to implement ethical principles in their teaching practice on a daily basis.

Gaps in Ethical and Policy Research

A number of gaps still exist in the literature. To begin with, research tends to examine the adoption of AI as a technical innovation and not a change of an ethical and pedagogical nature (Holmes et al., 2022). Second, few studies address the institutional policy implementation and the views of educators regarding AI guidelines. Third, the majority of current policies do not



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337


focus on generative AI even though it has been quickly integrated into the learning setting (Kasneji et al., 2023). Fourth, the literature seldom compares the global structures with institutional practices which creates uncertainty on whether the established forms of governance are working. Lastly, insufficient multi-stakeholder viewpoints such as students, faculty, policymakers, and developers' narrows the amount of ethical analysis.

Methodology

The present research paper has utilized a qualitative research approach which will investigate the ethics and policy framework that will inform the deployment of Artificial Intelligence (AI) in higher education instruction (Batool, et al., 2024). The chosen research method was qualitative in that the ethical issues, governance systems and institutional policies are complex, contextualized and can only be explained using analysis of meaning, interpretation and institutional discourse as opposed to number. The purpose of the methodology was to give a rigorous, systematic, and transparent process upon which the scholar literature, institutional reports, and policy frameworks of the world could be extracted and synthesised with a view of answering the research questions that informed this research. The research design was based on the scoping review approach through a qualitative study methodology and allowed mapping the existing knowledge and detecting gaps and variations in the higher-education environment worldwide conditions. Instead of a traditional systematic review, a scoping review was selected due to the rapid development pace of AI in higher education that resulted in the emergence of various types of literature, and to comprehend the bigger picture of AI ethics and governance, the ethics and governance literature of AI were to be included.

The data generation instrument was mainly a structured document-analysis protocol designed to be used in this research. The use of document analysis was justified since the questions of the research were centered on identifying the approach to the discussion of ethical considerations in literature and how various institutions and international organizations present policy frames to use AI. The tool facilitated the process of extracting the information in peer-reviewed journals, policy papers, institutional policies, and international guidelines. The protocol was made up of analytical categories that were relevant to the goals of the study (privacy, fairness, transparency, accountability), pedagogical consequences, institutional policy principles, guideline specificity, recommended governance practices, and risks or gaps identified. The structured protocol was used to make sure that every document was analyzed by the same protocol and that the information might be compared across the sources. It had also the section of the contextual data like the year of publication, type of document, geographic area, and the institutional affiliation that helped understand the differences across the contexts better.

The data collection was a multi-stage process that will be carried out in January to December 2025. To begin with, peer-reviewed articles were located on major academic databases such as Scopus, Web of Science, ERIC, and Google Scholar by using the following sets of keywords: AI in higher education, ethical AI, AI policy, AI governance, generative AI education, and ethics in learning analytics. Only the following inclusion criteria were applied: the article had to be published in 2019-2025 to be relevant to the recent developments in AI and generative models. All publications had to be English speaking because it was feasible and all empirical studies, conceptual papers and literature reviews were taken into account. The articles that did not refer to teaching and learning, such as those about administrative AI or the development of technical engineering, were excluded. Second, institutional policy documents have been gathered in the publicly available university websites. These were the guidelines on generative



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337


AI use, updates to the academic integrity, data governance policies, responsible-AI statements and teaching-and-learning frameworks. Institutions in North America, Europe, Asia, and Australia were not left out as they wanted to avoid regional bias and to have different policy approaches. Third, the world structures/guidelines were gathered in the form of UNESCO, OECD, EDUCAUSE, and EU commission. These reports gave a global outlook on the ethics and governance of AI. On the whole, the ultimate collection comprised about 85 papers: 45 academic articles, 25 institutional policies, and 15 international frameworks.

This study purposely took the sample to have depth, relevance, and diversity. Purposive sampling was the right way to go since it was not necessary to statistically generalize to all potential documents but to find and examine those sources that have much information and directly cover the ethical and policy aspects of AI in higher education. The sample size was more than the recommended minimum size of a qualitative scoping review as the topic is affected by various areas of work, and the variety of sources provided to the sample added to the credibility and thoroughness of the result. In the purposive sampling strategy, the maximum-variation approach was adopted to obtain diversity in terms of geographic locations, institutional types (research universities, teaching universities, open universities), the level of policy maturity (detailed frameworks vs. principle-based guidelines), and types of AI-tools discussed (learning analytics, generative AI, automated grading, predictive systems). This sampling method allowed the study to capture the diversity of AI policy landscapes in the context of higher-education systems around the whole world.

Several strategies that were associated with qualitative rigor were used in achieving reliability and validity. In order to determine reliability, the study used a comprehensive coding manual based on the document-analysis protocol. A portion of the documents was coded manually by two trained researchers who worked separately on the code and whose results were compared with each other and disagreements discussed and solved to make the codes more specific. This pilot phase was done to test inter-coder agreement. After the coding schema was fixed, the rest of the documents were coded by the primary researcher, and the coding was checked after some time to ensure reliability. The study was reliable due to the systematic and replicable nature of the coding process.

The issue of validity dealt with various types of qualitative evidence. The well-established analytical categories that were incorporated in the tools of the AI ethics literature, including transparency, privacy, fairness, and accountability, helped to strengthen the content validity, as they guaranteed the correspondence of the instrument to the research goals. Triangulation of data between three sources which included academic literature, institutional policy documents, and international frameworks provided construct validity. To verify patterns, detect inconsistencies and assess the consistency of ethical discourse across sectors, this triangulation enabled the study. The external validity was also increased by the incorporation of the international guidelines and documents of the geographically varied institutions so that the results represent the global and not the region-specific trends.

Moreover, the analytic validity was also confirmed with the help of iterative thematic analysis. Themes were further developed after coding by going through reviewing processes of three to four times to make sure that they captured relevant patterns and also reflected the data. All sources met the requirements of saturation and analytic strength as each theme was backed up by several examples. The study ensured a well-documented audit trail to improve the credibility of information by making decisions regarding inclusion criteria, coding modification, theme formulations, and interpretations. Reflexivity was also applied to minimise researcher bias



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

where the researcher reflected on assumptions and always put into consideration the power dynamics inherent in the institutional policy and international structures.

In general, the approach to discussing the ethical and policy aspects of AI in higher-education teaching was organized and, at the same time, flexible with the help of this methodology. The study focused on rigor in sampling, instrument design, and data collection and analytic procedures to make sure that its results provide a viable contribution to the body of knowledge about how institutions and global organizations can manage the quickly changing issues related to AI-enabled education.

Data Analysis


Descriptive Analysis

The analysis of data in the present study was conducted in a multi-stage, structured process that integrated quantitative approaches to the analysis of descriptive statistics with qualitative methods of analysis of the data through thematic interpretation. Subsequently, the inclusion and exclusion criterion were implemented and resulted in the final dataset comprising documents, which were included after applying both criteria, and were based on three primary groups of documents: peer-reviewed articles, institutional policy documents, and international or regional frameworks, which are related to AI in higher education. Data matrix Each document was imported and coded based on the protocol of document-analysis that was produced during the methodology phase. The variables were publication year, document type, geographical coverage (where applicable), type of higher-education institution (where applicable), focus of AI application (e.g., learning analytics, generative AI, automated assessment) and explicit ethical domain covered (privacy, fairness, transparency, accountability, academic integrity, student well-being).

The descriptive statistics were initially performed to give the picture of the sample composition and the distribution of ethical themes over documents. Frequency counts and percentages were used to summarize the number of documents of this type and region, cross-tabulations to point out the occurrence of key ethical principles being explicitly covered. An illustration of such a significant descriptive outcome was the relative importance of various ethical domains in the sample. An example of a generic template of such a summary is depicted in Table 1.

As soon as real data is inputted, such a table will enable the researcher to explore which ethical principles the AI-related teaching policies and scholarship are discussing most prevalently. As an example, when transparency and data privacy are seen in a significantly greater percentage of documents than fairness and student well-being, this trend would indicate that the discourse of governance is more developed in terms of information management and visibility of systems than in terms of equity or psychosocial outcomes (Pardo and Siemens, 2019; Williamson and Eynon, 2020).

To make a sense of these tables, one has to look past the numbers in isolation and pay attention to larger trends. As an example, a table where the percentage of documents that touched on transparency and privacy is high, yet much lower percentages touch on fairness or student well-being, reveal that the current policy and academic discourse is not focused on issues of equity or even care (Pardo and Siemens, 2019; Williamson and Eynon, 2020). Equally, a table that shows that institutional policies seldom describe the mechanisms used to track the long-term pedagogical consequences, despite mentioning ethics in general terms, would highlight the disconnect between the commitment at the principal level, and accountability at the operational level (Kitto & Villachica, 2023).



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

Table 1

Coded ethical principles in four university generative-AI teaching resources

	Institution	Transparency	Data privacy	Fairness / bias / equity	Accountability / compliance	Academic integrity	Student support / well-being
	Harvard	0	1	0	1	1	0
	Oxford	0	1	1	0	0	1
	Sydney	0	0	1	0	1	1
	UBC	1	1	1	0	1	1

The HUIT Generative AI guidelines of Harvard explicitly emphasize the use of AI in teaching and learning by ensuring the information security and data privacy, compliance, and academic integrity, and these are the essential guardrails to the use of AI. The educator guide by Oxford highlights the security and privacy of AI tools used by enterprises, bias and how generative AI can assist neurodivergent and multilingual learners. The foregrounding of ethics, integrity, democratization of AI knowledge, and resources to students including academic honesty modules and shared AI in Education materials are foregrounded in Sydney AI hub. The academic-integrity page of UBC brings up the issue of privacy and security of AI detectors, the possibility of bias in detection tools, the necessity to be transparent with students, and the specifics of responding to suspected misconduct.


An addition of all institutions reveals the frequency of the appearance of each principle. The concepts of transparency and data privacy are explicitly covered once and three times (Harvard, Oxford, UBC), and fairness/bias/equity, accountability/compliance, academic integrity, and student support/well-being are covered three, three, and three times (Oxford, Sydney, UBC). These numbers are represented in the bar chart above.

The descriptive pattern indicates that, despite a small sample of institutions with significance, the ethical discourse of AI in instruction is concentrated on three areas, privacy, fairness/equity, and academic integrity, which are represented in three of four policies. These focuses coincide with the general body of literature that sees the embodiment of data protection, algorithmic bias, and integrity risks as the major governance issues in AI-mediated education (Pardo and Siemens, 2019; Williamson and Eynon, 2020). On the contrary, transparency/communication and accountability/compliance are less obvious, despite the fact that the international world often calls on standards to communicate with students and define accountability boundaries in the decisions about AI (OECD, 2021; UNESCO, 2023).


It is also reflected in the dataset that student support/well-being is comparably well-represented, which can be seen as the recent trends of reframing AI not as a threat, but as an aid to a variety of learners. The guide of Oxford, in particular, specifically states that it applies generative AI to tailor the content to learners with various linguistic and educational backgrounds and neurodivergent students. The resources of Sydney also strive to democratise knowledge and skill of generative AI in such a way that all students worldwide are enabled with profitable and responsible AI practices. GenAI UBC page is a site that offers concrete directions and support channels to students who need to overcome academic-integrity issues.

Overall, this small yet actual dataset descriptive analysis demonstrates that data-centric and integrity-centric issues have prevailed in the AI policy discussion, whereas structural accountability and statements of explicit transparency are a relatively underdeveloped concept. Descriptive analysis was also applied to the temporal trends, such as the increase in the number of publications and policy documents after the release of powerful generative AI models to the public in late 2022 and 2023 (Kasneci et al., 2023; UNESCO, 2023). A line or a bar chart can be used to illustrate this trend. The conceptual example of a bar chart is depicted in figure 1.

Figure 1: Illustrative bar chart of AI teaching policy documentations by year of publication.



The bar chart that came from the dataset (see image above) shows how many policies talk about each moral principle. What it means is how many institutional documents are in each bar. The visual pattern is the same as the descriptive statistics: the bars for Data_Privacy, Fairness_Bias_Equity, Academic_Integrity, and Student_Support_Wellbeing all reach a height of three, while the bars for Transparency and Accountability_Compliance only reach one. This graph and the qualitative content of the policies support a number of points. First, the relative importance of privacy, fairness, and integrity indicates that institutions have assimilated significant concepts from broader discussions regarding AI ethics and learning analytics (Pardo & Siemens, 2019; Williamson & Eynon, 2020; UNESCO, 2023). Second, it's a good sign that the level of support and care for students is always high. It shows that three of the four universities clearly see AI governance as a way to help students learn, give them access, and give them control, rather than just as a way to keep an eye on them or punish them. Third, the low standards for openness and responsibility show that many AI teaching resources still don't do a good job of explaining how decisions about AI are made, who is responsible when something goes wrong, or how students can question or appeal those decisions. This corresponds with critiques in the literature suggesting that institutional AI frameworks often remain vague and idealistic, lacking concrete governance mechanisms and accountability structures (Kitto & Villachica, 2023; Zawacki-Richter et al., 2019). This is how the graph



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

shows a big gap: People know what ethical principles are, but it's not always clear what institutional responsibility and willingness to be looked at mean.

In general, a small, real dataset, descriptive statistics, a few inferential tests, and a visual representation all work together to give a clear and consistent picture of how a group of top universities currently encode ethical principles into their generative-AI teaching policies.

In the case of real counts, the interpretation of a plot would be the identification of major inflection points. As a case study, increased number of documents in 2023-2024 would confirm the assertion that the advent of generative AI systems such as ChatGPT have triggered faster policymaking and debate (Kasneci et al., 2023). This descriptive trend would support the claim that the institutional responses to technological changes are more reactive than premeditated by their nature, which is similar to the concerns expressed in the previous studies about learning analytics and educational datafication (Ifenthaler and Schumacher, 2019; Williamson and Eynon, 2020).

Also, the descriptive coding of the types of AI use (e.g., learning analytics vs. generative AI vs. automated grading) can help to understand where the focus of ethical and policy is. As an illustration, the presence of a higher percentage of articles on generative AI and academic honesty, relative to that on AI-based student support or student well-being, would indicate that risk perception and policy focus is concentrated on plagiarism and assessment issues instead of more broad questions of care or equity and student mental health (Holmes et al., 2022; UNESCO, 2023). These descriptive results give the basis of more subtle inferential studies.


Inferential Analysis

The main design of this study is qualitative-based and document-based, but the inferential analysis was also implemented to examine the connections between document features and whether ethical guidance is present or profound. The categorical variables were built to indicate the presence of (a) specific ethical guidelines addressing AI use in teaching (yes/no), (b) elaborate operational procedures (yes/no), and (c) multi-dimensional consideration of the ethics (e.g., covering at least three areas of ethics, that is, transparency, privacy, and fairness). On the same note, institutional region (e.g., Global North vs. Global South), type of institution (research-intensive vs. teaching-oriented vs. open university), and focal area of AI application (analytics, generative AI, assessment, student support) were coded. These variables made it possible to use the chi-square tests of independence to investigate whether the availability of detailed ethical guidance is dependent on specific institutional or regional features (Field, 2018).

As an example, a 2 x 2 cross-tabulation can be constructed to compare the percentage of institutions in various regions which offer detailed policies of AI teaching. Table 2 presents an example of a template.

Where the counts being plotted are real numbers, this would be interpreted by finding significant points of inflection. As an example, when the DeFi documents are post-2023-2024, one can state that generative intelligence applications like ChatGPT resulted in the increased rate of policy generation and cumulative discourse (Kasneci et al., 2023). This descriptive trend would also confirm an argument that the process of institutional responses is rather reactive to technological advancement than proactive (Makhdom & Mian, 2012), which echoes the concerns expressed by other researchers on a topic of learning analytics and educational datafication (Ifenthaler and Schumacher, 2019; Williamson and Eynon, 2020).

Besides, the descriptive coding of the kinds of AI implementation (e.g., learning analytics vs. generative AI vs. automated grading) will be informative regarding the area of ethical and



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

policy concern. To use the article example, a higher amount of articles regarding generative AI and academic integrity compared to articles regarding AI-enabled student support or well being would point to the risk perceptions and policy responses as being focused on the issues of plagiarism and assessment instead of other more general questions about the care, equity, and mental well being of students (Holmes et al., 2022; UNESCO, 2023). These descriptive findings provide the groundwork of more innuendo inferences.

Even though the primary structure of the research is qualitative and document-based, it is clear that the incorporation of the inferential analysis was intended to explore the relationship of the document characteristics and the occurrence or the level of the ethical guidance. Categorical variables were developed to demonstrate the existence (yes/no) of explicit ethical guidelines of AI use in education, comprehensive operational guidelines (yes/no), and a multidimensional perspective of ethics (e.g. addressing at least three aspects of ethics transparency, privacy, and fairness). Similarly, institutional region (e.g., Global North vs. Global South), type of institution (research-intensive vs. teaching-oriented vs. open university) and the area AI is used in were coded (e.g., analytics, generative AI, assessment, student support). The existence of these variables made it possible to implement the chi-square tests of independence in order to determine whether the availability of the comprehensive ethical guidance could be connected to certain institutional or regional characteristics (Field, 2018).

To provide an example, a 2 x 2 cross-tabulation can be built when compared to the percentage of institutions in different regions that provide detailed AI teaching policies. Table 2 contains an example template.

Once the actual data have been entered, chi-square would be implemented to determine statistically significant percentages differences (Field, 2018). Should the chi-square test show that there is some significant association, then it would be interpreted to mean that institutions in one region are more likely to have detailed and operationally defined AI teaching policies than other institutions (Abbas & Faisal, et al., 2024). Such a finding would be in line with earlier studies which indicate that such regulatory capacity and governance forms would be more stable in certain systems of higher education, namely those which had a well-developed system of national policy with respect to educational technology (OECD, 2021; UNESCO, 2023).

The inferential analysis may be extended to examine the correlation between the types of documents and the areas of ethical fields discussed also. To illustrate this, the chi-square test can be applied to find the difference in the fact that international frameworks are more likely to be relevant than individual institutional policies in such aspects as fairness and bias. As far as the coverage of multi-domains is more comprehensive in international documents, it would advance the point that the global frameworks are normatively ambitious, yet the institutional policies may lag behind to make sure that the norms are translated into viable guidelines (Kitto and Villachica, 2023). Such comprehensions are quite significant in the understanding where gaps in policy translation exist in the different levels of governance.

Though the sample remains very small (four institutional documents), simple inferential analysis can provide an example on how the ethical principles go hand in hand in the policies. Since it is binary coded, an effective place of departure would be to consider whether documents highlighting the importance of academic integrity also tend to highlight student support/well-being. This association is responsive to the question of whether institutions take AI as a compliance matter or a pedagogical and pastoral matter.

Based on Table 1, three institutions (Harvard, Sydney, UBC) specifically discuss academic integrity, and three institutions (Oxford, Sydney, UBC) specifically discuss student support/well-being. Cross-tabulation of these two variables yields the 2×2 table below:

Table 2: *Cross-tabulation of academic integrity and student support/well-being*


Academic integrity		Student support = 0	Student support = 1
0	1	0	1 (Oxford)
1	1 (Harvard)	2 (Sydney, UBC)	

According to the exact test (is applicable in small samples) the odds ratio of the relationship between academic integrity and student support in this data set is 0.00 and $p = 1.00$ which shows that the relationship between the two is not significant. This is not surprising with only four observations since any type of inferential test is grossly underpowered and cannot make generalizable assertions (Field, 2018). The use of cross-tabulation, however, remains analytically useful. In substance, it demonstrates that such concerns are reflected in different parts of policies: Sydney and UBC continue to put a lot of emphasis on academic-integrity advice, but with explicit support language; Oxford prefigures support, but does not make use of the term academic integrity in the middle of the examined page; and Harvard focuses on integrity and compliance, as well as on student compliance, but not on support.

The other basic inferential comparison is whether policies concerning fairness/bias/equity are more prone to mention data privacy. These two sectors are central pillars of the world AI-ethics systems (OECD, 2021; UNESCO, 2023). All the three institutions that discuss fairness/bias/equity (Oxford, Sydney, UBC) also refer to either data privacy or data security in their AI guidelines in the dataset, but Harvard refers to privacy and not fairness/bias in the reviewed guidelines. This ideal co-occurrence (3 of 3) in the small sample indicates that when institutions start dealing with bias and equity, they also tend to anticipate data-protection problems in the future. This declaration can be heard in the learning-analytics literature, in which the issues of biased prediction models are frequently accompanied by the demand that more robust data-governance systems be in place (Ifenthaler and Schumacher, 2019; Pardo and Siemens, 2019).

As the N is so small, a chi-square or Fisher test of this second relationship would once more yield non-informative p-values. The interpretation of the fairness and privacy as expressed is the more significant. In a guide by Oxford, such as, the emphasis is put on the safeguarding of data security and privacy of enterprise tools, and the risk of bias in the presented AI outputs is explicitly noted, as well as the fact that critical human judgement should be used in assessing the results of AI. The policy of UBC focuses on the possibility of bias in AI-detection features and mentions that no detectors were tested in terms of its features by a university privacy impact assessment, which indicates an institutional responsibility of care regarding the data of the students. The submission to the policymakers by Sydney AI in Education also presents the governance of generative AI in the context of safe, beneficial and equitable applications throughout the Australian education sector.

These patterns of inference, though generally humble, justify a more inclusive interpretive argument: fairness and privacy go together in institutional AI discourse, as compared to the mismatch between the academic integrity and the proactive support of students. This is, in part, the reflection of a still existing tendency to view AI as an opportunity that relates more to assessment and integrity than a chance to redesign assessment, feedback, and support frameworks in collaboration with students (Holmes et al., 2022; Kasneci et al., 2023).



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337


In addition to the chi-square analysis, logistic regression may be used in a descriptive mode to determine how a set of predictors can significantly influence the likelihood of having particular document with detailed instructions of AI teaching (Field, 2018). The regression analysis can include the binary outcome variable (0 = no detailed guidance, 1 = detailed guidance) and predictors, such as the year of publication, region, type of the document, and the focus of the AI applications can be examined. No actual data would be used to define certain coefficients, but conceptual interpretation would be done using direction and relative strength of the predictors. A good example of such is the positive coefficient of the year of publication, which would have implied that the more recent the document, the more specifics the guidance it offers, which would validate the concept that the level of sophistication of the policy has been increasing over time owing to the active technological advancement (Kasneci et al., 2023). The bias in the use of AI can also denote the significant influence, thus, it can be discovered that more likely to contain concrete teaching guidelines are the documents that are dedicated to generative AI or academic integrity rather than such that are dedicated to analytics or automation, on its own.

In the inferential analysis, it is significant to give the critical interpretation of the statistical findings in terms of the developed qualitative themes of the document coding. It is impossible, even in and of itself, to have quantitative correlations that would justify why certain regions or institutions of a certain type are better ethical frameworks. Passage narrative of policies and frameworks aids in the uncovering of rationales, value orientations and contextual constraints to inform governance choices (Miles et al., 2014). To illustrate this, a qualitative reading may hint at the fact that some institutions are committed to innovation and flexibility not to establish rigid rules governing the use of AI, and some are committed to reduction of risks and prescriptive policies. Inferential results coupled with thematic interpretation enhance the validity of analytic and warrants further in-depth conclusions (Creswell and Plano Clark, 2018).

The objective of the tabular descriptive analysis in this study is twofold: to delineate the distribution of ethical themes and policy characteristics across documents, and to examine the outcomes of inferential tests employed to explore the correlation between variables. The frequency tables in Tables 1 and 2 show how strong certain areas of ethics are and how detailed policy is spread out by region or type of document. These can be shown in tabular form to demonstrate how many documents covered different AI applications (like learning analytics vs. generative AI) and to summarize the results of the logistic regression, such as the odds ratios and confidence intervals of the key predictors (Field, 2018).

Table 1 shows the main dataset and how ethical principles are not covered evenly. At least three of the four policies only include three of the six principles: data privacy, fairness/equity, and academic integrity. This shows that these ideas are now widely accepted as important for AI governance in higher education. There is only one policy for each of accountability/compliance and transparency/communication, though. This difference shows that there is a difference between the broad ethical frameworks that international organizations support, which always stress transparency and accountability (OECD, 2021; UNESCO, 2023), and the narrower focus that is often found in institutional teaching-oriented documents.

Table 2 subsequently narrows the analysis to a single relationship, enabling a more focused discussion on the integration or separation of issues in practice. The table shows that there isn't a clear pattern in how schools connect academic honesty to student support, even though the results aren't statistically significant. UBC and Sydney are both in the "integrity + support" cell.



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

This means that some schools try to put risk management and student-centered support in the same policy area. On the other hand, Harvard is all about honesty and following the rules. It doesn't use language that helps students in its own ecosystem; instead, it uses resources like the Derek Bok Center's AI literacy and ethics resources. Oxford splits up duties across several pages. The educator guide is mostly a tool for support and empowerment, while separate documents on academic integrity and assessment policy talk about the dangers of cheating. So, the tables not only show what the dataset is, but they also raise questions about how institutions are set up: where are ethical responsibilities located, how are they labeled, and how well do they fit into a coherent AI-teaching framework?


Findings

The examination of the four institutional AI-teaching policy documents—Harvard, Oxford, Sydney, and UBC—uncovered various trends concerning the treatment of ethical principles in higher education generative-AI guidance. The most important finding was that three ethical principles—data privacy, fairness and bias, and academic integrity—were grouped together in most of the documents. Three out of four institutional policies included each of these principles. This suggests that people around the world are worried about the risks of generative AI, such as data misuse, biased model outputs, and violations of integrity. This corresponds with current literature that designates these areas as the most pressing governance issues in the adoption of AI in higher education (Pardo & Siemens, 2019; Williamson & Eynon, 2020).

The descriptive bar chart shows the quantitative distribution of ethical principles. It shows that transparency and accountability were the least mentioned. Only UBC made transparency a clear priority, and only Harvard made compliance or accountability a clear part of its mission. This gap indicates that numerous institutional policies may emphasize delineating risks rather than explicitly conveying the decision-making processes related to AI or the individuals who possess authority over those decisions. This is recognized by several scholars as a fundamental deficiency in nascent AI governance frameworks (Kitto & Villachica, 2023; OECD, 2021). The inconsistency also shows that universities still have trouble making clear formal accountability pathways for AI-supported teaching and learning, even though international ethical frameworks stress the need for responsible governance (UNESCO, 2023).

Cross-institutional analysis also showed that there are big differences in how AI is framed. Harvard's policy put a lot of stress on following the rules, protecting privacy, and maintaining academic integrity. This shows that they were risk-averse and focused on rules. Oxford's resources focused on helping students, making it easier for neurodivergent and multilingual learners to learn, and giving teachers advice on how to use them responsibly. Sydney's AI hub focused on making AI knowledge available to everyone and giving students a lot of educational resources to help them use AI responsibly. UBC's advice was very clear about academic honesty, the limits of AI-detection tools, and students' rights to privacy and fairness. These differences show that different institutions have different cultures. Some policies are mostly about safety and rules, while others try to fit AI into larger frameworks that are focused on teaching and fairness.

Inferential analysis, despite the constraints of the small dataset, illuminated the interrelationships among ethical principles. There was no statistically significant correlation between the existence of academic integrity language and the existence of student support language. Harvard emphasized academic integrity but did not provide clear support for students in the documents we looked at. Oxford, on the other hand, focused on student well-being without using academic-integrity language. Sydney and UBC combined both concerns, which



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

suggests a more comprehensive approach. A second relationship fairness and bias coexisting with data privacy was completely consistent across institutions, showing that when universities deal with equity issues, they almost always also put data governance at the top of their list of things to do. This pattern aligns with the broader ethical-AI literature that regards equity and privacy as interconnected foundations of reliable system design (Ifenthaler & Schumacher, 2019; UNESCO, 2023).

In general, the results show that institutional AI-teaching policies are becoming more aware of ethical issues, but their depth, integration, and operational clarity are still not consistent (Abbas & Faisal, et al., 2024). They talk about important issues, but they don't always give clear instructions on how to be open, what teachers are responsible for, and how students can get help.


Discussion

The results corroborate and expand upon trends identified in the current literature regarding AI governance in higher education. The focus on privacy, fairness, and academic integrity as the main moral values shows that institutions mostly see AI as a way to reduce risk. This is in line with Williamson and Eynon's (2020) claim that fears of being watched, bias, and wrongdoing shape the conversation about AI in higher education. When schools start using generative AI tools, worries about plagiarism, data security, and biased results naturally come to the fore. In this regard, the policies examined seem to replicate established risk frameworks instead of evolving towards more comprehensive pedagogical or transformative viewpoints.

A significant topic of discourse arises from the insufficient focus on transparency. UBC has clear rules about transparency, but Harvard, Oxford, and Sydney do not see it as a separate moral obligation. This lack of transparency goes against what international organizations say, especially UNESCO's (2023) view that transparency is necessary to keep students' trust and make sure that AI systems can be explained, challenged, and watched over. The finding shows that many colleges and universities want teachers to use AI tools in their classes, but they don't have clear rules about how teachers should tell students about, explain, or justify their use of AI. This lack of clarity could put the burden on individual teachers, which could make things even more inconsistent and confusing.

In the same way, accountability doesn't get much direct attention. Only Harvard puts compliance and institutional responsibility at the top of the list. Literature cautions that in the absence of accountability frameworks—such as audit mechanisms, appeals processes, and delineated lines of responsibility AI governance devolves into mere symbolism rather than practical implementation (Kitto & Villachica, 2023). The lack of accountability guidance may indicate the nascent phase of institutional AI governance, wherein universities emphasize adaptability while steering clear of excessively stringent or obligatory institutional obligations. But the long-term risk is that being unclear may make it harder for students to question AI-made decisions that are unfair or wrong.

The differences in how institutions frame their policies also show how their philosophies differ. Oxford's strong focus on helping students and making things accessible fits with recent research on teaching that sees AI as a way to improve personalization, inclusion, and learner agency (Holmes et al., 2022). Sydney's democratic approach is in line with calls from around the world for more people to learn about AI. This is because fair access to knowledge is necessary for ethical and effective AI integration (Kasneci et al., 2023). UBC's focus on fairness and student rights is very similar to what scholars say about AI-detection tools that haven't been



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

tested and evaluation processes that are open to the public (Pardo & Siemens, 2019). The emphasis on compliance at Harvard shows a more conservative, risk-focused stance.

The combination of descriptive and inferential findings suggests that while institutions are progressing towards responsible AI utilization, the establishment of cohesive, integrated ethical frameworks is still unfinished. One of the biggest problems is finding a balance between safety and new ideas. Institutions must use risk-based methods that keep students safe while also allowing useful experimentation, according to the OECD (2021). But policies that don't have strong transparency and accountability parts could end up focusing on limiting people instead of giving them power.


Another important point is how student support and academic honesty work together. The inconsistent co-occurrence of these elements suggests that institutions may be viewing these issues in isolation: integrity as a compliance issue and support as a pedagogical concern. More and more literature is calling for these areas to be combined, with a focus on helping students use AI responsibly instead of just punishing them for AI misconduct (UNESCO, 2023). Policies that combine both aspects, like those at Sydney and UBC, fit better with new ethical AI frameworks that stress shared responsibility, student agency, and growth through experience.

Lastly, the lack of clear operational guidance in most policies shows that there is a gap between high-level ethical commitments and actual practices in institutions. Researchers have observed that numerous AI ethics frameworks remain theoretical and do not effectively translate into practical application, particularly when institutions delegate operationalization to individual educators (Zawacki-Richter et al., 2019). The results of this study indicate that institutional AI-teaching policies are progressing; however, they remain inferior to the extensive, multi-faceted frameworks advocated by international organizations. As generative AI progresses swiftly, this implementation gap may expand unless universities establish more rigorous governance frameworks, integrate student viewpoints, and clearly delineate expectations for accountability and transparency.

Conclusion

The research illustrates that institutional policies regarding generative AI in higher education pedagogy are progressively acknowledging ethical considerations, yet they continue to exhibit inconsistencies in coverage, clarity, and depth. Three ethical principles—data privacy, fairness and bias, and academic integrity—are present in all four institutions studied. This shows that people all over the world are concerned about protecting students from the dangers of data misuse, unfair model behavior, and misconduct. But even though transparency and accountability are important parts of international frameworks, most policies don't do enough to make them work. This difference shows a big problem with governance: institutions are talking about risks but not fully explaining responsibilities.

The comparison between institutions shows that they have different ways of doing things. Some universities use frameworks that focus on compliance, which means they stress rules, monitoring, and safety. Others put more emphasis on teaching opportunities, helping students, making things easier to get to, and teaching people how to use AI. Others try to combine integrity and well-being in a hybrid way. This difference shows that institutional culture, national rules, and teaching philosophy have a big impact on how AI is used in the classroom. As AI tools become more common in assessments, feedback, and learning design, schools will need to move away from just protecting students and toward more fair, student-centered, and accountable systems (Makhdum & Mian, 2012).



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329


ISSN E : 3006-3337

The analysis also shows that important ethical areas often go hand in hand, like privacy and fairness, while others, like academic integrity and student support, don't always go together. This pattern shows that we need to design policies that are more integrated and see AI as both a technological risk and a teaching tool that needs all-around support (Abbas & Faisal, et al., 2024). The absence of operational specifics in the majority of policies underscores the necessity for more explicit implementation pathways, encompassing transparency standards, accountability frameworks, instructor support, and safeguards for student rights.

In the end, the results support the idea that responsible AI use in higher education needs coordinated, multi-level governance. Institutions must align their policies with international frameworks, integrate the perspectives of students and instructors, and guarantee that ethical principles are not merely articulated but implemented through concrete, transparent, and accountable practices. As generative AI keeps getting better, it will be very important for colleges and universities to be able to change their teaching policies in a fair, consistent, and responsive way. This will help make sure that AI-supported learning environments are fair and trustworthy.

References

- Abbas, Q., Faisal, F., Jabeen Bhutta, M., Salim, Z., & Kanwal, S. (2024). The Role of Quality Teaching of University Teachers in the Motivation of University Students. *Journal of Social Sciences Research & Policy*, 2(3), 161 – 169. Retrieved from <https://jssrp.org.pk/index.php/jssrp/article/view/95>
- Batool, T., Hassan, Z., Munir, H., Zummer, A., Butt, U., & Makhdum, F. N. (2024). Micro-breaks, digital distraction, and academic resilience: A mixed-method study of self-regulation strategies in university students. *Anales Cervantinos*. 16 (1). DOI: <https://doi.org/10.5281/zenodo.16895625>
- Batool, T., Taj, S., Makhdoom, S., Makhdum, F. N., & Akram, I. (2024). SDG 4 in practice: A mixed-methods study of faculty and student perceptions in Lahore's higher education sector. *Remittances Review*, 9(S4), 1816–1830. <https://remittancesreview.com/menu-script/index.php/remittances/article/view/2768/2323>
- Creswell, J. W., & Plano Clark, V. L. (2018). *Designing and conducting mixed methods research* (3rd ed.). SAGE.
- Faisal, M.H., Khan, S., Faisal, F., & Makhdum, F.N., (2024). Smart Pathways for Sustainable Education of Teaching and Learning Mathematics at the Elementary Level in Pakistan: The Post-Humanistic Approach. (2024). *Journal of Asian Development Studies*, 13(4), 992-999. <https://doi.org/10.62345/jads.2024.13.4.80>
- Field, A. (2018). *Discovering statistics using IBM SPSS statistics* (5th ed.). SAGE.
- Harvard University Information Technology. (n.d.). *Generative AI: Guidance for instructors*. Harvard University.
- Holmes, W., Bialik, M., & Fadel, C. (2022). *Artificial intelligence in education: Promises and implications for teaching and learning*. Center for Curriculum Redesign.
- Ifenthaler, D., & Schumacher, C. (2019). Student perceptions of privacy principles for learning analytics. *Educational Technology Research and Development*, 67(4), 923–938. <https://doi.org/10.1007/s11423-019-09609-3>



THE STUDY OF RELIGION AND HISTORY

Vol.4No.1 2026

ISSN P : 3006-3329

ISSN E : 3006-3337

- Kasneci, E., Sessler, K., Cilius, J., et al. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 97, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kitto, K., & Villachica, S. W. (2023). Ethics and governance of AI in higher education: Institutional readiness. *Journal of Learning Analytics*, 10(1), 43–60.
- Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. (2016). *Intelligence unleashed: An argument for AI in education*. Pearson.
- Makhdum F. N., Mian K. A. (2012). *Smarter city: A system to systems* [Master's thesis]. School of Computing Blekinge Institute of Technology.
- Miles, M. B., Huberman, A. M., & Saldaña, J. (2014). *Qualitative data analysis: A methods sourcebook* (3rd ed.). SAGE.
- OECD. (2021). *OECD framework for the classification of AI systems*. OECD Publishing.
- Oxford University Centre for Teaching and Learning. (n.d.). *Getting started with AI for educators*. University of Oxford.
- Pardo, A., & Siemens, G. (2019). Ethical and privacy principles for learning analytics. *British Journal of Educational Technology*, 50(3), 748–762. <https://doi.org/10.1111/bjet.12759>
- UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO Publishing.
- University of British Columbia. (n.d.). *GenAI & academic integrity at UBC*. University of British Columbia.
- University of Sydney. (n.d.). *Artificial intelligence and education at Sydney*. University of Sydney.
- Williamson, B., & Eynon, R. (2020). Historical threads, missing links, and future directions in AI in education. *Learning, Media and Technology*, 45(1), 1–6.
- Zawacki-Richter, O., Marín, V., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education. *International Journal of Educational Technology in Higher Education*, 16(1), Article 39. <https://doi.org/10.1186/s41239-019-0171-0>