UNIVERSITY GOVERNANCE OF GENERATIVE ARTIFICIAL INTELLIGENCE: A COMPARATIVE ANALYSIS OF GUIDELINES IN GLOBAL ELITE UNIVERSITIES

CAMILA LOPES FERREIRA

Departamento de Engenharia de Produção, Universidade Tecnológica Federal do Paraná, Brazil
clferreira@utfpr.edu.br | https://orcid.org/0000-0001-6314-8764

ALESSANDRA DUTRA

Departamento de Ciências Humanas, Universidade Tecnológica Federal do Paraná, Brazil
alessandradutra@utfpr.edu.br | https://orcid.org/0000-0001-5119-3752

SANI DE CARVALHO RUTZ DA SILVA

Departamento Académico de Matemática, Programa de Pós-Graduação em Ensino de Ciência e Tecnologia, Universidade Tecnológica Federal do Paraná, Brazil
sani@utfpr.edu.br | https://orcid.org/0000-0002-1548-5739

JOSÉ ROBERTO HERRERA CANTORANI

Departamento de Educação Física, Instituto Federal de Educação, Ciência e Tecnologia de São Paulo, Brazil
cantorani@ifsp.edu.br | https://orcid.org/0000-0003-1792-0383

CLAUDIA TANIA PICININ

Departamento de Engenharia de Produção, Universidade Tecnológica Federal do Paraná, Brazil
claudiapicinin@utfpr.edu.br | https://orcid.org/0000-0003-4844-3516

ABSTRACT

This study comparatively analyses the institutional guidelines of ten global elite universities, selected based on the Times Higher Education (THE) World University Rankings 2025, regarding the use of generative artificial intelligence (AI). It adopts a qualitative, descriptive-analytical approach, grounded in documentary analysis of 87 institutional documents organised into four axes: teaching, student use, information security, and governance/ethics. Each axis was coded according to a three-level robustness scale, allowing the identification of patterns, gaps, and innovations in institutional responses. The results indicate the consolidation of a normative-formative model that articulates principles of algorithmic justice, transparency, and accountability, repositioning universities as key actors in the ethical and strategic framing of generative AI. The study contributes to the emerging field of university governance of AI by offering a critical systematisation of current practices and guidance for the development of institutional policies.

KEY WORDS

generative artificial intelligence; university governance; institutional policies; higher education; digital ethics.

SISYPHUS

JOURNAL OF EDUCATION

VOLUME 14, ISSUE 02,

2026, PP 248-279

DOI: https://doi.org/10.25749/sis.44224

GOVERNANÇA UNIVERSITÁRIA DA INTELIGÊNCIA ARTIFICIAL GENERATIVA: ANÁLISE COMPARATIVA DE DIRETRIZES EM UNIVERSIDADES DE ELITE GLOBAL

CAMILA LOPES FERREIRA

Departamento de Engenharia de Produção, Universidade Tecnológica Federal do Paraná, Brasil
clferreira@utfpr.edu.br | https://orcid.org/0000-0001-6314-8764

ALESSANDRA DUTRA

Departamento de Ciências Humanas, Universidade Tecnológica Federal do Paraná, Brasil
alessandradutra@utfpr.edu.br | https://orcid.org/0000-0001-5119-3752

SANI DE CARVALHO RUTZ DA SILVA

Departamento Académico de Matemática, Programa de Pós-Graduação em Ensino de Ciência e Tecnologia, Universidade Tecnológica Federal do Paraná, Brasil
sani@utfpr.edu.br | https://orcid.org/0000-0002-1548-5739

JOSÉ ROBERTO HERRERA CANTORANI

Departamento de Educação Física, Instituto Federal de Educação, Ciência e Tecnologia de São Paulo, Brasil
cantorani@ifsp.edu.br | https://orcid.org/0000-0003-1792-0383

CLAUDIA TANIA PICININ

Departamento de Engenharia de Produção, Universidade Tecnológica Federal do Paraná, Brasil
claudiapicinin@utfpr.edu.br | https://orcid.org/0000-0003-4844-3516

RESUMO

Este estudo analisa comparativamente as diretrizes institucionais de dez universidades de elite, selecionadas com base no Times Higher Education (THE) World University Rankings 2025, quanto ao uso da inteligência artificial (IA) generativa. A investigação adota uma abordagem qualitativa, de natureza descritivo-analítica, fundamentada na análise documental sistemática de 87 documentos institucionais, organizados em quatro eixos: ensino, uso por estudantes, segurança da informação e governança/ética. Cada eixo foi codificado segundo uma escala de robustez em três níveis, o que permitiu identificar padrões, lacunas e inovações nas respostas institucionais. Os resultados indicam a consolidação de um modelo normativo-formativo que articula princípios internacionais, como justiça algorítmica, transparência e responsabilidade, e reposiciona as universidades como protagonistas na formulação ética e estratégica do uso da IA generativa. O estudo contribui para o campo da governança universitária da IA ao oferecer uma sistematização crítica das práticas vigentes e subsídios para instituições em processo de estruturação de suas políticas.

PALAVRAS-CHAVE

inteligência artificial generativa; governança universitária; políticas institucionais; ensino superior; ética digital.

SISYPHUS

JOURNAL OF EDUCATION

VOLUME 14, ISSUE 02,

2026, PP 248-279

DOI: https://doi.org/10.25749/sis.44224

GOBERNANZA UNIVERSITARIA DE LA INTELIGENCIA ARTIFICIAL GENERATIVA: ANÁLISIS COMPARATIVO DE DIRECTRICES EN UNIVERSIDADES DE ÉLITE GLOBAL

CAMILA LOPES FERREIRA

Departamento de Engenharia de Produção, Universidade Tecnológica Federal do Paraná, Brasil
clferreira@utfpr.edu.br | https://orcid.org/0000-0001-6314-8764

ALESSANDRA DUTRA

Departamento de Ciências Humanas, Universidade Tecnológica Federal do Paraná, Brasil
alessandradutra@utfpr.edu.br | https://orcid.org/0000-0001-5119-3752

SANI DE CARVALHO RUTZ DA SILVA

Departamento Académico de Matemática, Programa de Pós-Graduação em Ensino de Ciência e Tecnologia, Universidade Tecnológica Federal do Paraná, Brasil
sani@utfpr.edu.br | https://orcid.org/0000-0002-1548-5739

JOSÉ ROBERTO HERRERA CANTORANI

Departamento de Educação Física, Instituto Federal de Educação, Ciência e Tecnologia de São Paulo, Brasil
cantorani@ifsp.edu.br | https://orcid.org/0000-0003-1792-0383

CLAUDIA TANIA PICININ

Departamento de Engenharia de Produção, Universidade Tecnológica Federal do Paraná, Brasil
claudiapicinin@utfpr.edu.br | https://orcid.org/0000-0003-4844-3516

RESUMEN

Este estudio analiza comparativamente las directrices sobre el uso de la inteligencia artificial (IA) generativa en diez universidades de élite, seleccionadas con base en el Times Higher Education (THE) World University Rankings 2025. La investigación adopta un enfoque cualitativo, de carácter descriptivo-analítico, sustentado en el análisis sistemático de 87 documentos institucionales organizados en cuatro ejes: docencia, uso estudiantil, seguridad de la información y gobernanza/ética. Cada eje fue codificado según una escala de robustez de tres niveles, lo que permitió identificar patrones, brechas e innovaciones en las respuestas institucionales. Los resultados señalan la consolidación de un modelo normativo-formativo que articula principios internacionales, como justicia algorítmica, transparencia y responsabilidad, y reposiciona a las universidades como protagonistas en la formulación ética y estratégica del uso de la IA generativa.

PALABRAS CLAVE

inteligencia artificial generativa; gobernanza universitaria; políticas institucionales; educación superior; ética digital.

SISYPHUS

JOURNAL OF EDUCATION

VOLUME 14, ISSUE 02,

2026, PP 248-279

DOI: https://doi.org/10.25749/sis.44224

University Governance of Generative Artificial Intelligence: A Comparative Analysis of Guidelines in Global Elite Universities

Camila Lopes Ferreira[1], Alessandra Dutra, Sani de Carvalho Rutz da Silva, José Roberto Herrera Cantorani, Claudia Tania Picinin

INTRODUCTION

One of the most consistent ways of anticipating educational trends and transformations is to observe the practices and guidelines adopted by institutions that occupy leading positions in the global arena. Universities located at the frontier of knowledge and academic innovation tend to respond more quickly to emerging challenges, such as the advance of artificial intelligence (AI), especially in its generative form. Analysing what these universities are doing, at the institutional level, in the face of the expansion of AI in higher education makes it possible not only to map patterns of action, but also to identify potential pathways for other institutions that seek to adapt to this new context.

Generative AI has emerged as one of the most disruptive technologies of the early twenty-first century, rapidly transforming communicational practices, knowledge production processes, and teaching and learning structures. In the educational field, AI has facilitated the collaborative production of texts and arguments, enriching communication (Cress & Kimmerle, 2023), reducing cognitive load, and promoting the collective construction of knowledge through the automated generation of texts, images, videos, and data analyses (Chen & Wu, 2024; Mittal et al., 2024).

However, its expansion has brought not only pedagogical innovations but also deep epistemological and ethical questions that challenge traditional paradigms of authorship, assessment, and teaching mediation. As language models such as ChatGPT, Bard, and Claude become widely accessible tools, universities worldwide are compelled to revise their academic policies to address the impacts of generative AI on scientific integrity, equity in access to information, and the security of institutional data. This movement forms part of broader transformations of the university as a formal organization embedded in a globalized society, marked by regimes of evaluation, rankings, and accountability devices that reconfigure the principles, purposes, and modes of regulation of higher education (Cowen, 2013; Ramirez, 2013).

Despite the relevance of the topic, the literature reveals a significant gap in the comparative systematization of institutional guidelines adopted by universities regarding the use of generative AI. Most studies focus on general normative analyses or isolated case studies, without providing a consolidated overview of how internationally renowned institutions have addressed this issue. This gap is particularly notable at a time when international bodies, such as the United Nations Educational, Scientific, and Cultural Organization (UNESCO, 2022) and the European Commission (2021), have called on educational institutions to play a leading role in developing ethical principles and governance strategies for AI. Furthermore, the absence of qualitative metrics to assess the

[1] Universidade Tecnológica Federal do Paraná, Rua Doutor Washington Subtil Chueire, número 330, Jardim Carvalho, Ponta Grossa, Paraná, Brasil.

degree of robustness of these guidelines limits the capacity of other universities, especially in developing countries, to adopt models that are both inspiring and contextually adapted.

In this study, institutional responses to generative AI are interpreted through the concept of a normative formative model of university governance. This model refers to the articulation between, on the one hand, normative devices, such as regulations, codes of conduct, information security protocols, and academic integrity policies, and, on the other hand, formative devices, such as AI literacy programmes, pedagogical guides, professional development initiatives, and ethical awareness activities. Unlike strictly regulatory approaches focused solely on control and sanction, the normative formative model emphasises the role of the university as an organization that, while regulating the use of AI, simultaneously educates and produces new frames of meaning regarding technology, authorship, and responsibility in higher education.

LITERATURE REVIEW AND CONCEPTUAL FRAMEWORK

To strengthen the consistency between the literature review and the comparative model employed in this study, we ground the normative formative approach in two complementary contributions that specify central dimensions for institutional governance of generative AI. Huttenlocher et al. (2023) frame AI governance as a purposeful institutional endeavour aimed at enabling beneficial uses while addressing key societal and organisational priorities—such as security, individual privacy and autonomy, safety, shared prosperity, and democratic and civic values. In this perspective, governance is operationalised not only through principles but also through institutional mechanisms that support accountability and public understanding of system performance, including disclosure practices and audit-oriented approaches (Huttenlocher et al., 2023). In higher education, these dimensions translate directly into the need for formal guidance on risk management, data protection and security constraints, as well as transparent institutional arrangements for responsible use.

In turn, Porsdam Mann et al. (2024) offer an ethical framework for the use of large language models in academic writing that highlights three cumulative requirements: (i) human responsibility for accuracy and integrity through meaningful human review; (ii) substantial human intellectual contribution; and (iii) acknowledgement and transparency regarding the role of LLMs in producing academic outputs. While developed for research and scholarly communication, these requirements clarify the ethical baseline that universities increasingly convert into teaching and assessment guidance and into student-facing rules on acceptable use, disclosure, and accountability (Porsdam Mann et al., 2024).

Taken together, these contributions provide the conceptual dimensions that make the normative formative model analytically tractable at the institutional level. In this study, they are operationalised through four thematic axes—teaching, student use, information security, and governance/ethics—and through a three-level ordinal robustness scale that captures a progression from (1) indicative or generic mentions without detail or institutionalization, to (2) formal guidelines with restricted scope or only incipient actions, and finally to (3) consolidated guidelines articulated with institutional actions and accompanied by evidence of practical implementation, as reflected in official university documents (Huttenlocher et al., 2023; Porsdam Mann et al., 2024).

Beyond these scholarly contributions, recent policy-oriented frameworks and recommendation documents converge around four expectations that are especially

relevant to higher education: transparency in AI-assisted work, human accountability, protection of personal and research data, and institutional oversight able to monitor risks while enabling innovation (European Commission, 2021; Russell Group, 2023; UNESCO, 2022; United States, Office of Science and Technology Policy, 2022). These documents do not replace empirical institutional analysis, but they provide an important normative background against which university guidelines can be interpreted.

Taken together, the literature and policy frameworks indicate that the governance of generative AI in higher education is no longer limited to classroom experimentation; it now involves institutional capacity to formalise expectations, allocate responsibilities, and translate broad ethical principles into operational rules. Yet comparative studies examining how universities publicly codify these expectations across teaching, student use, information security, and governance structures remain scarce. This study positions itself in that gap by analysing the degree of robustness with which such dimensions are formalised in leading universities.

Against this backdrop, this study aims to compare and analyse the institutional guidelines of global elite universities regarding the use of generative AI, with an emphasis on pedagogical, student-related, information security, and ethical governance dimensions. The sample comprises ten higher education institutions with the best performance in the Times Higher Education (THE) World University Rankings 2025 (Times Higher Education, 2025). The systematic documentary analysis is based on a qualitative analytical matrix constructed from four thematic axes and three levels of robustness. By offering an integrated view of the most advanced educational policies in the field of generative AI in high-performing universities, this study seeks to contribute to the formulation of good practices and critical frameworks that can guide other institutions in the development of responsible, inclusive, and effective strategies for the regulation and use of AI in higher education, as well as to deepen the theoretical debate on university governance in algorithmically mediated contexts.

METHOD

This study adopts a qualitative, descriptive-analytical design grounded in systematic documentary analysis of institutional policies regarding generative artificial intelligence (GenAI) in global elite universities. The methodological protocol follows established principles for documentary research, including the selection of relevant institutional records, screening by explicit inclusion criteria, and iterative interpretation of texts in relation to an analytical framework (Bowen, 2009). The unit of analysis is the publicly available institutional document published in official university communication channels (e.g., policy statements, teaching and assessment guidance, academic integrity policies, IT/security guidelines, ethics/governance reports).

CRITERIA FOR SELECTING INSTITUTIONS

Universities were selected based on the Times Higher Education (THE) World University Rankings 2025, focusing on the ten institutions with the highest overall performance in the global ranking (Times Higher Education, 2025). THE was chosen due to its

multidimensional approach combining indicators related to teaching, research, citation impact, knowledge transfer, and internationalisation, which makes it a relevant reference for identifying institutions with strong academic influence and policy-setting capacity. This sampling strategy privileges universities in high-income contexts with established trajectories in adopting educational technologies, and therefore emphasises highly visible cases with substantial institutional capacity to produce formal guidelines.

DATA COLLECTION PROCEDURES

Data collection occurred between January and April 2025, using exclusively publicly available primary sources from the official websites of the ten selected universities. The corpus included documents such as: teaching and assessment guidance for staff; academic integrity and authorship statements; rules and recommendations regarding students' use of GenAI tools; information security and data protection guidance (including approved/licensed platforms); and institutional documents addressing AI ethics and governance.

Document inclusion followed three criteria:

- *Currency:* publication date or last update compatible with the period under analysis, ensuring documents reflected guidelines in force up to April 2025.
- *Institutional authorship*: explicit attribution to the university or a clearly identified official unit (e.g., central administration, teaching and learning office, IT/security office).
- *Topical relevance*: direct applicability to at least one of the four analytical axes (teaching; student use; information security; governance/ethics).

All eligible materials were catalogued in a structured database (document title, issuing unit, date/update, URL, and preliminary thematic tagging). In total, 87 documents were analysed and organised by institution and by analytical axis.

ANALYTICAL AXES, CONCEPTUAL DIMENSIONS, AND CODING PROCEDURE

Conceptual anchoring and dimensions of analysis

The analysis model was defined a priori based on specialised literature addressing governance and responsible deployment of GenAI, with particular attention to: (i) transparency and accountability requirements in academic contexts; and (ii) risk- and responsibility-oriented governance approaches relevant to institutional settings (Huttenlocher et al., 2023; Porsdam Mann et al., 2024).

Two conceptual inputs are central for the operationalisation adopted in this study:

- *Accountability, transparency, and responsible use in academic work*: Porsdam Mann et al. (2024) emphasise human responsibility for accuracy and integrity, substantial human contribution, and acknowledgement/transparency regarding LLM use. These dimensions informed the Student use axis (e.g., disclosure rules, authorship integrity, accountability) and the Teaching focus axis (e.g., assessment guidance and how institutions frame permissible pedagogical uses while preserving academic integrity).
- *Governance through risk-aware oversight and institutional mechanisms*: Huttenlocher et al. (2023) highlight the need for governance mechanisms that help users understand system capabilities/limits and that support oversight through institutional arrangements (e.g., disclosure practices, auditing/assurance logic, and priority concerns such as security and privacy). These dimensions informed the Information security axis (e.g., treatment of sensitive data, approved tools, security constraints) and the Governance and ethics axis (e.g., institutional structures, ethical principles, training/oversight arrangements).

Together, these conceptual dimensions were translated into four analytical axes that capture how universities regulate GenAI in teaching and learning, how they guide students' use, how they manage information risks, and how they institutionalise governance and ethical oversight.

Analytical axes and operational categories

The documentary analysis was organised around four thematic axes:

- *Teaching focus*: pedagogical guidance and didactic strategies for using GenAI in teaching and assessment (e.g., assessment redesign, AI literacy, permissible instructional uses, guidance for instructors).
- *Student use*: institutional orientations on acceptable student use, authorship and integrity, disclosure, accountability, and (when present) equity considerations related to access to tools.
- *Information security*: rules and recommendations for handling sensitive data, approved/licensed platforms, restrictions on public tools, and alignment with data protection and security requirements.
- *Governance and ethics*: internal frameworks for AI oversight, ethics principles, transparency and fairness considerations, training initiatives, and institutional structures supporting responsible GenAI deployment.

Documents were multi-coded when they addressed more than one axis (e.g., a teaching guidance document that also included disclosure requirements and security cautions).

Coding rubric (three-level ordinal scale) with operational examples

To enhance transparency and replicability, coding followed a three-level ordinal rubric applied per axis:

- Level 1 — Basic: the axis is mentioned only in indicative or generic terms, with limited operational detail and no clear institutionalisation (e.g., broad statements encouraging "responsible use" without specifying rules, processes, or concrete practices).
- Level 2 — Intermediate: formal guidance exists, but it is limited in scope, fragmented across units, or focused primarily on cautionary recommendations with limited procedural specification (e.g., course- or discipline-level rules without broader institutional framing; partial guidance that lacks clear processes for disclosure, training, or enforcement; security cautions without an explicit approved-tool policy or risk procedure).
- Level 3 — Robust: consolidated guidance is identifiable, articulated with institutional actions and accompanied by evidence of implementation mechanisms (e.g., centrally issued instructor guidance and training resources; explicit student disclosure/accountability rules; security protocols specifying handling of sensitive data and/or approved tools; formal governance/ethics structures, frameworks, and institutional initiatives).

Because the study relies on documentary evidence, "implementation" was operationalised conservatively as the presence of institutional mechanisms that indicate enactment beyond a statement of principle—such as formal training materials, structured guidance for staff, approved-tool policies, established centres/committees/initiatives, or clearly defined institutional procedures (as reflected in the documents analysed). Illustratively, within the corpus: instructor-facing guides and AI literacy resources exemplify robust Teaching focus; explicit disclosure expectations and integrity framing exemplify robust Student use; approved/licensed tool guidance and restrictions on sensitive data exemplify robust Information security; and identifiable institutional structures and initiatives related to ethics and oversight exemplify robust Governance and ethics.

From document coding to institutional scores (scoring matrix)

Coding proceeded in two linked steps:

- Step 1 — Document-level coding: each document received (a) axis tags (one or more axes) and (b) a provisional level (1–3) for each axis it substantively addressed, based on the rubric above. Document coding captured what the text concretely established (rules, processes, responsibilities, security constraints, training/oversight mechanisms), not merely the presence of general claims.
- Step 2 — Institution-level scoring by axis: for each university and axis, the full set of coded documents was synthesised into a single institutional score

(1–3). This score represents the overall robustness of the institution's publicly stated guidance for that axis, considering: (i) existence and formality of guidance; (ii) specificity and scope (how detailed and how institution-wide vs. limited/fragmented); and (iii) institutional mechanisms indicating enactment (e.g., training resources, approved-tool policies, governance structures).

Where an institution presented both centrally issued guidance and more localised documents, scoring prioritised the most institution-wide and formally endorsed position evidenced in the corpus, while remaining consistent with the rubric's distinction between restricted scope (Level 2) and consolidated institutionalisation (Level 3). The resulting institution-by-axis scoring matrix is reported in the comparative results table to allow readers to trace the analytical pathway from coding decisions to cross-institutional comparison.

To make the analytical trajectory more explicit, the sample codes in Table 1 illustrate how recurrent documentary patterns were translated into provisional document-level classifications. The examples are synthetic excerpts derived from patterns observed across the corpus and are presented only to clarify the coding logic.

Table 1

Robustness scores (1–3) of institutional guidelines on AI in the ten institutions with the best performance in THE 2025, by thematic axis

Documentary Cue / Excerpt	**Analytical Decision and Rationale**
"Use AI tools responsibly and verify generated content before submission".	Axis: student use. provisional code: level 1 (basic). Rationale: the text signals responsible use, but does not specify disclosure procedures, limits of acceptability, or institutional mechanisms.
"In this course, GenAI may be used only when authorised by the instructor; any use must be acknowledged in the submitted assignment".	Axis: student use / teaching focus. provisional code: level 2 (intermediate). Rationale: there is explicit guidance, but its scope is localised and dependent on course-level discretion rather than institution-wide policy.
"Do not upload confidential, personal, or research-sensitive data to public AI tools; use only institutionally approved platforms. AI-assisted work must be disclosed according to university guidance".	Axis: information security / student use. provisional code: level 3 (robust). Rationale: the document defines clear restrictions, specifies approved tools, and links operational rules to institutional procedures.

Coder independence and adjudication

Two researchers independently performed the coding. After independent coding, axis-level classifications and provisional scores were compared. Disagreements were discussed and resolved through consensus, with reference to the rubric definitions, to strengthen consistency and methodological rigour. This adjudication step also served to refine the application of categories and ensure stable interpretation of borderline cases across institutions and axes.

DATA ANALYSIS AND VISUALISATION STRATEGY

The analysis combined within-case synthesis and cross-case comparison. Within each institution, coded evidence was summarised by axis to identify the dominant regulatory approach and the presence of institutional mechanisms (e.g., training and guidance structures, security constraints, governance initiatives). Across institutions, the consolidated axis scores were compared using the scoring matrix, enabling pattern identification (convergences, gaps, and divergences) and supporting the formulation of recommendations grounded in observed institutional practices.

Scores were then visualised in a comparative figure designed for interpretability with ordinal, multi-institution data, allowing readers to quickly identify differences by axis and by institution while preserving the qualitative interpretation underpinning the scoring.

RESULTS

This section presents the results of the documentary analysis conducted on institutional guidelines regarding the use of AI in internationally renowned universities. The information was organised in two articulated stages: a detailed description of each institution's policies, based on available official documents, and a comparative synthesis of the findings structured according to the previously defined thematic axes.

This strategy aims to provide a descriptive and analytical perspective on the current scenario, enabling the identification of patterns, specificities, and gaps in institutional approaches to AI.

DESCRIPTION OF INSTITUTIONAL GUIDELINES BY THE UNIVERSITY

This subsection presents, in a systematic way, the guidelines of each of the ten higher education institutions with the best performance in the Times Higher Education (THE) World University Rankings 2025 regarding the use of AI. The analyses are organised in accordance with the official documents of each university, highlighting actions related to the axes of teaching, student use, information security, and governance/ethics.

University of Oxford

The University of Oxford has adopted a comprehensive and responsible approach to the use of AI, with guidelines that span teaching, information security, institutional communication, and research ethics. These orientations seek to align the use of AI with principles of academic integrity, institutional responsibility, and critical innovation.

In teaching and assessment, the university recognises the potential of generative AI as a support for learning, provided that it is used according to ethical and pedagogical criteria (University of Oxford, 2025a). It adopts the principles of the Russell Group (2023),

which emphasise the promotion of AI literacy among students and staff, the training of teachers to guide effective use of these tools, the ethical and equitable adaptation of teaching and assessment methods, the preservation of academic integrity, and the promotion of the sharing of good practices. The university values didactic strategies that develop critical thinking, authorship, and metacognition, especially in writing and problem-solving activities. Unauthorised use of AI in assessments is classified as academic misconduct (University of Oxford, 2025a).

Students are provided with specific guidance on the safe and responsible use of generative AI, which must be acknowledged in their work where its use is permitted. The university also recommends that students follow any additional guidance provided by their departments or courses, according to disciplinary specificities (University of Oxford, 2025e).

In institutional communication, the University of Oxford warns of the risks associated with automatically reproducing AI-generated content, emphasizing the need for critical review and alignment with institutional values (University of Oxford, 2025b). In terms of information security, the use of licensed platforms, such as Microsoft 365 Copilot and ChatGPT Edu, is recommended, provided that a prior risk analysis is conducted and sensitive data is protected (University of Oxford, 2025d).

In the ethical domain, the Institute for Ethics in AI promotes interdisciplinary debates on the moral and political implications of technology (University of Oxford, 2025c). The university also contributed to the development of guidelines on the use of large language models (LLMs) in scientific writing, published in the journal *Nature Machine Intelligence*, which highlight transparency, substantial human authorship, and critical review as fundamental principles (Porsdam Mann et al., 2024). Taken together, these actions demonstrate Oxford's institutional commitment to the responsible use of AI, grounded in ethics, security, and academic excellence.

Massachusetts Institute of Technology (MIT)

The Massachusetts Institute of Technology (MIT) adopts a comprehensive and continually updated institutional approach to the ethical, safe, and responsible use of AI, especially in the context of teaching, research, and university management. Its guidelines reflect the institution's commitment to technological innovation associated with principles of responsibility, transparency, and equity.

In the field of information security, the Department of Information Systems and Technology (IS&T) recommends the preferential use of licensed tools such as Microsoft Copilot to ensure compliance with legislation such as the *Family Educational Rights and Privacy Act* (FERPA), which protects the privacy of students' education records (U.S. Department of Education, 2011), the *Health Insurance Portability and Accountability Act* (HIPAA), which establishes federal standards for the privacy and security of health information (U.S. Department of Health & Human Services, 2026), and Massachusetts data protection regulations, including 201 CMR 17.00 and M.G.L. c. 93H, which establish standards for safeguarding personal information of Massachusetts residents (Massachusetts, 2025a, 2025b). These tools should not be used to process confidential data, make critical decisions, or handle sensitive intellectual property without a prior formal risk analysis. Additionally, the importance of transparency in the use of AI tools in academic and administrative activities is emphasized, with a focus on traceability and the accuracy of the generated information (MIT Information Systems and Technology, 2025).

In teaching, MIT Sloan guides its faculty through the document 'Getting started with AI-enhanced teaching: a practical guide for instructors', which presents theoretical foundations, use cases, recommended tools, institutional support, and ethical guidelines (MIT Management, 2025). The material proposes the use of AI for creating examples, quizzes, and interactive strategies, always with critical review, data protection, and bias mitigation in mind. It is also recommended that the use of AI by students remain optional, thereby promoting equity and intellectual autonomy.

MIT also contributes to the public debate on AI governance policies through the Schwarzman College of Computing and the MIT Washington Office, advocating regulation oriented towards the purpose of use (Huttenlocher et al., 2023). The MIT Media Lab, in partnership with Harvard's Berkman Klein Center, leads the Ethics and Governance of Artificial Intelligence project, which promotes a public interest-oriented approach (MIT School of Architecture + Planning, 2023).

The MIT Internet Policy Research Initiative (2024) examines policies related to reliability, transparency, and privacy to enhance trust in social applications of AI. In the ethical education of young people, institutions develop curricula for basic education that focus on algorithmic justice, critical thinking, and the social implications of AI (Payne, 2019). Taken together, these actions underscore MIT's leadership role in promoting the ethical, democratic, and safe use of AI within and beyond the university environment.

Harvard University

Harvard University adopts a broad and integrated approach to the use of AI, with guidelines that link academic integrity, information security, research ethics, and institutional governance. The official orientations address faculty, students, and administrators, covering both benefits and risks associated with tools such as ChatGPT and other generative AI platforms, with recent updates that reflect developments in 2025.

The Office of the Provost has issued guidelines that establish clear parameters for the use of generative AI, including the prohibition on entering confidential or non-public research data, in accordance with the university's Information Security Policy. The guidelines emphasize the importance of transparency and critical review of AI-generated content, given the potential for inaccuracies, distortions, or copyright infringements. Academic units are encouraged to develop their own guidelines that align with institutional principles (Harvard University Office of the Provost, 2025).

In teaching, instructors are advised to communicate the rules governing students' use of AI clearly. Schools and departments may adopt approaches ranging from total prohibition to conditional permission, and it is essential that teachers explicitly state their expectations (Harvard University Office of the Provost, 2025).

Harvard Medical School (HMS) has established specific rules that prioritize the critical verification of AI-generated outputs, scientific integrity, and the protection of sensitive data. The use of tools such as Harvard AI Sandbox and ChatGPT Edu requires prior training and adherence to formal request protocols (Harvard Medical School, 2025).

Harvard University Information Technology (HUIT) maintains a portal that provides detailed guidance on licensed tools, permitted data levels, and good use practices, with an emphasis on privacy and reliability (Harvard University Information Technology, 2025).

The ethical dimension is strongly articulated by the Berkman Klein Center for Internet & Society, which leads the Ethics and Governance of Artificial Intelligence initiative in partnership with MIT Media Lab. This initiative promotes research on

algorithmic justice, explainability, social inclusion, and unequal impacts of automated technologies (Berkman Klein Center for Internet & Society, 2025).

Additionally, the Harvard Division of Continuing Education offers the course "AI Ethics in Business: Managing Bias and Ethical Usage," designed for organizational leaders seeking to mitigate bias and enhance ethical integrity in business practices (Harvard University, 2025). Through these initiatives, Harvard reaffirms its leading role in shaping responsible AI policies that combine innovation with ethical vigilance, institutional security, and civic education.

Princeton University

Princeton University adopts a comprehensive and responsible approach to the use of AI, with guidelines that address teaching, information security, institutional communication, and research ethics. The aim is to ensure that the use of technologies, such as generative AI, aligns with the principles of academic integrity, institutional responsibility, and critical innovation.

In teaching and assessment, Princeton acknowledges that tools such as ChatGPT can support learning, provided they are used in accordance with ethical and pedagogical criteria. The university advises faculty to establish clear policies on the use of AI in their courses, to encourage transparency on the part of students, and to preserve academic rigour. Unauthorised use of these tools in assessments may be characterised as academic misconduct, undermining the development of intellectual autonomy (Princeton University, 2025b).

For students, ethical and responsible use of AI is recommended, with the tools understood as support for knowledge construction rather than substitutes for academic authorship. Given the possibility of inaccuracies and biases, AI-generated content should be acknowledged in assignments where its use is permitted. Departments and schools are encouraged to develop specific guidelines tailored to the particularities of each field (Princeton University, 2025b; Princeton University Library, 2025).

In institutional communication, the university promotes the strategic use of generative AI while emphasizing the need for review, validation, and transparency, particularly when sensitive data or potentially misleading information is involved (Princeton University, 2025c).

Regarding information security, Princeton recommends the use of licensed tools such as Microsoft 365 Copilot, with institutional access. It stresses that sensitive data and intellectual property should not be entered into public platforms (Princeton University, 2025d, 2025e).

The ethical dimension is addressed through the Princeton Dialogues on AI and Ethics initiative, a collaboration between the University Center for Human Values (UCHV) and the Center for Information Technology Policy (CITP). The project aims to develop frameworks for critical reflection and guidance on responsible decision-making regarding AI, involving scholars from the humanities, social sciences, and exact sciences (Princeton University, 2025f). Additionally, the university hosts interdisciplinary events that explore issues such as privacy, justice, and responsibility (Princeton University, 2025a). These initiatives demonstrate Princeton's commitment to ethical, safe, and academically rigorous uses of AI.

University of Cambridge

The University of Cambridge has developed comprehensive and up-to-date guidelines on the use of AI, focusing on ethical responsibility, information security, pedagogical practices, and institutional governance. The guidelines are addressed to the entire academic community, including students, faculty, and administrative staff, to consolidate a culture of critical, safe, and ethical AI use.

In teaching and assessment, Cambridge recognises the potential of generative AI as a learning support, while warning of the risks of its indiscriminate application. The university recommends that departments and faculty establish clear policies based on AI literacy, equity of access, and academic integrity (University of Cambridge, 2025b). The guide published by the Blended Learning Service presents strategies for adapting assessment tasks, promoting original authorship, and reducing dependence on AI-generated responses (University of Cambridge, 2025a).

For students, the guidelines emphasize that AI can assist with idea development, study planning, and language revision, but should not replace independent intellectual work. The university stresses that, where permitted, the use of AI must be acknowledged in submitted assignments and that each school or discipline may impose more restrictive rules (University of Cambridge, 2025b).

In institutional communication, Cambridge has issued guidance on the use of generative AI by communication and marketing teams. These tools may be used for drafting initial texts and supporting creativity, provided that humans review the final content. Use in sensitive situations, such as crisis response, is discouraged due to the risk of inaccuracies and misinformation (University of Cambridge, 2025c).

On the ethical front, the Trust & Technology Initiative leads the project 'AI: Futures and Responsibility', in partnership with the Leverhulme Centre for the Future of Intelligence (LCFI) and the Centre for the Study of Existential Risk (CSER). This initiative examines the social and moral impacts of AI. It proposes recommendations based on principles such as justice, explainability, and autonomy, involving researchers from the sciences, humanities, and social sciences (University of Cambridge, 2025d). These actions position the University of Cambridge as a benchmark for striking a balance between technological innovation, academic integrity, and the public good.

Stanford University

Stanford University adopts a transparent and responsible institutional approach to the use of AI, with guidelines that address teaching, research, information security, and academic integrity. The orientations are based on principles that promote technological innovation without neglecting fundamental ethical and social values.

In teaching and assessment, Stanford recognizes the potential of generative AI tools, such as ChatGPT, as a learning support, while warning of the risks associated with their misuse (Stanford University, 2025c). Faculty are advised to include explicit policies on AI use in their syllabi, with approaches ranging from strict prohibition to encouraged use, provided that human authorship is clearly attributed (Stanford University, 2025d). The aim is to preserve intellectual originality and promote critical and ethical thinking (Stanford University, 2025c).

To support faculty, the university has produced a guide titled 'Teaching in the AI Era', which offers suggestions for planning classes, assignments, and assessments adapted to the new technological context. The material proposes activities that encourage students to critically reflect on AI-generated outputs, revising them in light of their reliability and relevance (Stanford University, 2025e).

The AI Meets Education initiative at Stanford (AImES), coordinated by the Center for Teaching and Learning, fosters interdisciplinary discussions among faculty, researchers, and students to explore the impact of AI on higher education. The initiative fosters an institutional culture based on equity, responsibility, and transparency, promoting a growth mindset and active participation (Stanford University, 2025a).

In 2025, Stanford published a set of institutional principles for the use of AI in academic, administrative, and communication activities. These principles emphasize the centrality of human responsibility, the preservation of individual autonomy, and the continuous monitoring of AI impacts, especially in assessment, selection, and decision-making processes (Stanford University, 2025b).

At the global level, the university contributes to ethical and regulatory debates through the AI Index Report, published annually by the Stanford Institute for Human-Centered Artificial Intelligence (HAI). The 2025 edition presents data and recommendations on the responsible use of AI across various sectors, including education (Stanford HAI, 2025). These actions consolidate Stanford's leading role in building academic policies that integrate technological innovation and ethical principles.

California Institute of Technology (Caltech)

The California Institute of Technology (Caltech) adopts clear and ethical guidelines for the use of AI, with a focus on academic integrity, information security, and social responsibility. The guidelines encompass processes ranging from admissions to pedagogical practices and interdisciplinary research, reflecting the institution's commitment to an innovative culture guided by ethical principles.

In admissions, the university issued a 2025 policy specifically on the use of AI in application essays. The use of generative AI for drafting or structuring texts is discouraged and considered acceptable only for grammatical revision or as a support for brainstorming. Applicants are required to disclose any use of the technology in their application materials, underscoring the importance of authenticity and personal expression (California Institute of Technology, 2025f).

In teaching, the Center for Teaching, Learning, and Outreach (CTLO) provides resources that guide faculty in defining clear policies on AI use in the classroom. The institution recommends that teachers address the role of these tools critically in students' education, encouraging ethical and pedagogical reflection (California Institute of Technology, 2025a). The Division of the Humanities and Social Sciences, for example, prohibits the use of generative AI without explicit instructor permission, classifying unauthorised use as an academic violation (California Institute of Technology, 2025b).

Regarding information security, Caltech warns that public AI platforms do not have confidentiality agreements with the institution. It is recommended that sensitive, personal, or proprietary data not be entered into such tools, and that institutional

guidance be sought before using AI in academic or administrative contexts (California Institute of Technology, 2025c).

In the ethical and cultural domain, the Linde Center for Science, Society, and Policy (LCSSP) promotes debates on algorithmic justice and scientific responsibility, integrating science, policy, and society (California Institute of Technology, 2025d). The university also organizes cultural initiatives, such as the MACH 33 festival, which brings together science, art, and ethics to explore social representations of AI in the contemporary imagination (California Institute of Technology, 2025e). These actions indicate Caltech's commitment to a critical, transparent, and ethically grounded approach to AI in its institutional practices.

University of California, Berkeley

The University of California, Berkeley (UC Berkeley) adopts an ethical and comprehensive approach to the use of AI, with policies that cover teaching, information security, institutional communication, and interdisciplinary research. Its guidelines are designed to ensure that the use of these technologies aligns with the institutional values of integrity, responsibility, and critical innovation.

In teaching, the university recommends that instructors establish clear policies on the use of generative AI, weighing both benefits and risks. The Teaching and Learning division highlights the importance of adapting pedagogical activities to promote critical thinking and student authorship. It is recommended that syllabi include explicit clauses on AI use, encouraging a reflective and transparent approach consistent with the learning objectives (University of California, Berkeley, 2025d).

For students, the university permits the use of generative AI in tasks such as grammatical revision or idea organisation, provided that the final content results from the student's own work. Unauthorised use in coursework or assessments may constitute a violation of academic rules. Where permitted, the use of AI must always be explicitly acknowledged (University of California, Berkeley, 2025c).

In institutional communication, UC Berkeley recommends the careful use of AI by media and marketing professionals. It warns of the risks associated with disseminating sensitive or inaccurate content. It emphasizes that any AI-generated text must be critically reviewed by humans before publication (University of California, Berkeley, 2025a).

Regarding information security, the university advises that sensitive, personal, or copyright-protected data should not be entered into public AI tools. The recommendation is to prioritise institutionally licensed platforms and to conduct risk assessments before any institutional use of these technologies (University of California, Berkeley, 2025c).

The university also hosts the AI Policy Hub, an interdisciplinary centre dedicated to research on public policy and the social impacts of AI. The programme brings together experts from various fields to address issues such as algorithmic justice, technological regulation and governance, aiming to influence institutional practices and contribute to evidence-based policies (University of California, Berkeley, 2025b). These initiatives highlight UC Berkeley's commitment to responsible and critical AI use, integrating technological innovation, academic excellence, and ethical values.

Imperial College London

Imperial College London has established robust guidelines for the use of AI, seeking to integrate this technology in an ethical, critical, and safe manner into teaching, research, communication, and institutional management. The university recognises the potential of generative AI, such as ChatGPT, to support learning and academic work, but stresses that its use must be grounded in pedagogical and ethical principles.

In teaching, faculty are advised to formulate explicit policies regarding the use of AI in their courses. The university's Teaching Toolkit provides practical guidance on incorporating AI into educational processes fairly and effectively, promoting digital literacy and critical thinking. AI is expected to function as a support, not a substitute, for academic authorship, and each department is responsible for defining the most appropriate approach according to its disciplinary specificities (Imperial College London, 2025b).

For students, the university library has issued recommendations on the responsible use of AI, emphasizing that generated content must be thoroughly checked for accuracy and appropriateness, and that its use, when permitted, should be clearly acknowledged. The guidelines also warn of the risks of excessive dependence on these tools, which may compromise intellectual autonomy and violate academic integrity norms (Imperial College London, 2025a).

In institutional communication, Imperial College permits the use of generative AI for tasks such as drafting and idea generation, but requires rigorous human review before publication, particularly in sensitive contexts. Human editorial curation is considered essential to ensure institutional responsibility and informational accuracy (Imperial College London, 2025a).

Regarding information security, the university discourages the use of public AI tools to handle sensitive data or intellectual property. It recommends adopting licensed platforms such as Microsoft Copilot, as indicated in documents issued by the library and academic services (Imperial College London, 2025a).

Additionally, the university promotes ethical training through the AI Fundamentals program, which is aimed at researchers, managers, and policymakers. The initiative promotes interdisciplinary discussions on algorithmic transparency, justice, and AI regulation, with a focus on the social and institutional challenges of the digital age (Imperial College London, 2025c). These guidelines and actions position Imperial College London as a global benchmark for the responsible adoption of AI, combining scientific excellence, ethical commitment, and institutional innovation.

Yale University

Yale University promotes a responsible and ethical institutional approach to AI use, with guidelines that encompass teaching, institutional communication, information security, and interdisciplinary research. The orientations aim to ensure that AI is used in alignment with the university's values, including academic integrity, intellectual autonomy, and a commitment to the public good.

In teaching, Yale recommends that faculty establish specific policies on the use of generative AI, taking into account pedagogical objectives and potential risks. The Poorvu Center for Teaching and Learning offers detailed guidance for instructors, emphasizing the critical use of technology and cautioning that dependence on AI may undermine

active learning and student authorship. The centre suggests that assignments and assessments be redesigned so that they remain meaningful even in contexts where these tools are available (Yale, 2025a).

For students, the university recognises that AI can assist with tasks such as language revision, idea organisation, and brainstorming. However, its use must be authorised in advance by the instructor and acknowledged in assignments where applicable. Yale emphasises that responsibility for the final content lies solely with the student and warns of the risks of inaccuracies, biases, and artificial coherence in AI-generated responses (Yale, 2025b).

In institutional communication, the Office of the Provost has issued guiding principles for the use of AI in academic and administrative communications. The guidelines emphasize the importance of human oversight and the preservation of the institutional voice, with AI viewed as a support tool rather than a substitute for editorial judgment. Use in sensitive contexts or involving confidential data is discouraged (Yale University, 2023).

Regarding information security, Yale recommends the use of licensed platforms, such as the Clarity Platform, which provide greater institutional control and oversight. Public AI tools should not be used to process personal, academic, or intellectual property-protected data, and careful analysis is required before adopting new solutions (AI at Yale, 2025a).

The university also promotes interdisciplinary initiatives such as the AI at Yale Symposium, which brings together specialists from multiple fields to discuss algorithmic justice, institutional responsibility, and the social effects of AI (AI at Yale, 2025b). These initiatives demonstrate Yale's ongoing commitment to critical innovation and the responsible use of AI, while upholding its academic and institutional values.

COMPARATIVE SYNTHESIS OF INSTITUTIONAL GUIDELINES

Based on the previous analyses, this subsection presents a synthesis of the institutional guidelines organised by thematic axis. Table 2 reports the robustness scores (1–3) assigned to each institution across the four analytical axes defined in this study (teaching, student use, information security, and governance/ethics), while Table 3 summarises the predominant approaches identified in the documentary analysis.

The combination of qualitative description and comparative systematisation makes it possible to identify both structural convergences and significant variations in how these universities have responded to the challenges posed by generative AI.

Table 2

Robustness scores (1–3) of institutional guidelines on AI in the ten institutions with the best performance in THE 2025, by thematic axis

University	Teaching Focus	Student use	Information security	Governance and ethics
University of Oxford	3	3	3	3
MIT	3	3	2	3
Harvard University	3	3	3	3

University	Teaching Focus	Student use	Information security	Governance and ethics
Princeton University	3	3	3	3
University of Cambridge	3	3	3	3
Stanford University	3	3	3	3
Caltech	2	2	2	3
UC Berkeley	3	3	3	3
Imperial College London	3	3	3	3
Yale University	3	3	3	3

Note: 1 = basic; 2 = intermediate; 3 = robust.

Table 3

Institutional guidelines on AI in the ten institutions with the best performance in THE 2025, organised by thematic axis

University	Teaching focus	Student use	Information security	Governance and ethics
University of Oxford	AI literacy, authorship, and critical thinking, Russell Group principles	Use as support; mandatory acknowledgement; local guidelines	Licensed tools; prohibition of use with sensitive data	Institute for Ethics in AI; interdisciplinary ethical framework
MIT	Critical integration of AI: a practical guide for instructors	Optional use with transparency; intellectual autonomy	Prohibition of use with confidential data without risk analysis	Partnerships with Harvard, Media Lab, and initiatives on algorithmic justice
Harvard University	Promotion of authorship; mandatory statements on AI use	Mandatory acknowledgement; critical review of AI-generated content	Licensed tools; compliance with privacy policies	Berkman Klein Center; ethics course and institutional policies
Princeton University	Faculty autonomy; preservation of authorship; clear course-level guidelines	Support for reasoning; mandatory acknowledgement; specific guidelines	Microsoft 365 Copilot; restriction on the use of public AI	Dialogues on AI and Ethics: an interdisciplinary approach
University of Cambridge	Adaptation of tasks and assessments; original authorship	Use permitted with restrictions; variation by discipline	Prohibition of sensitive data in public AI; preferential use of licensed tools	Trust & Technology Initiative; focus on justice and explainability
Stanford University	Teaching in the AI era guide: critical reflection	Variable models of use; encouragement of critique and revision	Institutional principles and continuous monitoring	AI Index Report: principles of institutional responsibility

University	Teaching focus	Student use	Information security	Governance and ethics
Caltech	Policies by discipline; prohibition unless authorised	Permitted only for grammatical revision or brainstorming; disclosure required	Avoidance of public AI with sensitive data; recommended institutional consultation	LCSSP; MACH 33 festival on science, art, and AI
UC Berkeley	Adaptation of pedagogical activities; AI clauses in syllabi	Use permitted with acknowledgement; violations treated as misconduct	Risk assessment required; preference for licensed platforms	AI Policy Hub; focus on regulation and social impact
Imperial College London	Guided use in the classroom with a focus on digital literacy	Mandatory verification; warnings about risks and explicit acknowledgement	Use of public AI discouraged; preference for internal solutions	AI Fundamentals; training on algorithmic justice
Yale University	Focus on authorship and active learning; policies by discipline	Transparency required; use permitted with instructor authorisation	Licensed platforms such as Clarity; public AI is not recommended	AI at Yale: ethical governance and explainability

In addition to the qualitative description of the guidelines adopted by each university (Table 3) and the robustness scoring matrix (Table 1), a comparative chart (Figure 1) was produced to synthesise the degree of robustness of the institutional AI guidelines adopted by the ten institutions analysed, distributed across the thematic axes of teaching, student use, information security, and governance/ethics.

Figure 1

Comparison of universities by thematic axis

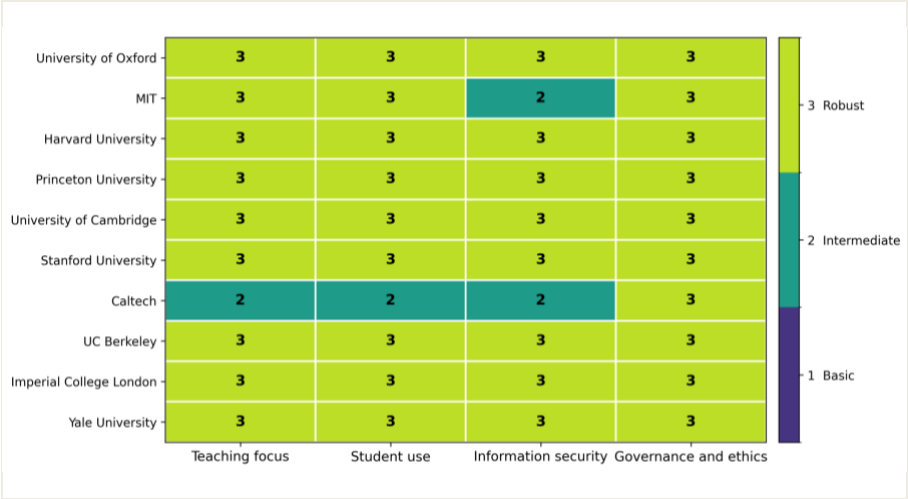

Note: 1 = basic; 2 = intermediate; 3 = robust.

Figure 1 visualises the robustness matrix reported in Table 2 and allows a rapid identification of where institutions differ by axis. In Teaching focus, nine universities score level 3 and one scores level 2 (Caltech), with no level-1 cases. The same distribution

is observed for Student use (nine at level 3; Caltech at level 2). Differences are slightly larger in Information security: eight universities score level 3, while MIT and Caltech score level 2. Governance and ethics is the only axis with full convergence at the top level, with all ten institutions coded at level 3.

These intermediate scores correspond to specific documentary features discussed in the institutional profiles. Caltech's level-2 scores in Teaching focus and Student use reflect guidance that is largely expressed through course- or division-level permissions and prohibitions (e.g., requirements for explicit instructor permission and restrictions in some units), rather than a consolidated, institution-wide pedagogical framework (Section 3.1.7). Caltech's level-2 score in Information security is supported by its warnings that public AI platforms operate without confidentiality agreements and by explicit recommendations not to enter sensitive, personal, or proprietary data and to seek institutional guidance before use (Section 3.1.7). MIT's level-2 score in Information security is grounded in IS&T guidance that prioritises licensed tools (e.g., Microsoft Copilot) and conditions the handling of confidential data, critical decisions, or sensitive intellectual property on prior formal risk analysis (Section 3.1.2).

Beyond these exceptions, the concentration of level-3 scores across axes indicates a shared institutional understanding of generative AI as a support technology that must be governed through a combination of pedagogical guidance, student-facing accountability and disclosure expectations, and safeguards for data and institutional responsibility. At the same time, the documents show that 'robust' (level-3) arrangements can follow different regulatory styles: some universities articulate common principles while delegating operational rules to courses or academic units (e.g., Harvard, Stanford, Yale, Cambridge), whereas others couple delegation with more centralised security guidance and approved-tool recommendations, as reflected in their institutional IT policies and portals.

Read together, Table 2, Table 3, and Figure 1 operate at complementary analytical scales: Table 3 preserves the documentary nuance of each institutional profile, Table 2 translates that evidence into an ordinal comparison, and Figure 1 makes cross-case concentration and dispersion immediately visible. This combined presentation prevents the visual synthesis from being interpreted in isolation. It clarifies that institutions with similar robustness scores may still differ in how they organise pedagogical guidance, disclosure rules, security constraints, and ethics structures.

DISCUSSION

The growing presence of generative AI in academic practices has prompted universities worldwide to develop specific policies regarding its use, particularly in the context of higher education. Against this backdrop, this study set out to analyse the institutional guidelines adopted by the ten higher education institutions with the best performance in the Times Higher Education (THE) World University Rankings 2025 (Times Higher Education, 2025), focusing on how generative AI has been incorporated into pedagogical, ethical, technological, and administrative dimensions.

The investigation employed a qualitative, descriptive-analytical approach, supported by systematic documentary analysis. A total of 87 institutional documents were examined and categorised according to four previously defined analytical axes: teaching, student use, information security, and governance/ethics. The analysis was conducted using a robustness scale and systematized through comparative visualizations, which

made it possible to identify patterns, gaps, and good practices adopted by internationally renowned institutions.

The results reveal not only the widespread adoption of regulatory policies concerning generative AI but also the emergence of a new lexicon of academic governance that repositions the relationships among technology, authorship, and integrity in higher education. The interpretation of the data reveals that, although there are variations in the emphasis and operationalization of guidelines, all the analyzed institutions converge in recognizing that generative AI cannot be separated from its epistemological, ethical, and social implications.

In the teaching axis, the incorporation of pedagogical practices that aim to utilize generative AI as a tool for cognitive development rather than mere productivity is particularly noteworthy. Several documents recommend the use of AI to foster critical thinking, argumentation, and ethical reflection (Gonsalves, 2024; Yan et al., 2024), explicitly linking technology to formative objectives. Among the universities analysed, nine obtained the maximum score (3), indicating consolidated guidelines with transversal coverage and strong articulation with the formative dimension. These institutions promote faculty digital literacy and curricular adaptation to algorithmically mediated contexts, often linking AI to continuous pedagogical development programmes. This approach aligns with UNESCO's (2022) recommendations, which advocate for strengthening critical thinking autonomy and preparing educators for the ethical and pedagogical challenges of AI. The fact that almost all universities have specific documentation aimed at teaching staff indicates a proactive approach to AI use, oriented towards critical mediation, creativity, and authorial integrity.

To keep pace with this incorporation, teaching staff are called upon to act as facilitators, experience designers, and co-creators of knowledge together with students and AI tools, which requires ongoing development and new digital competences (Kong & Yang, 2024; Mishra et al., 2024; Zhai, 2024). Only the California Institute of Technology (Caltech) received a score of 2 on this axis, with more restrictive policies and a narrower focus on encouraging critical reflection without a broadly formalised pedagogical orientation.

About student use, the findings indicate a consolidated level of normative maturity, with nine universities, including Harvard, MIT, Stanford, Oxford, and UC Berkeley, receiving the maximum score (3). These institutions have clear policies on limits of use, accountability mechanisms, and the promotion of transparency. Coding revealed that these universities not only define acceptable practices but also actively foster the development of an ethical digital culture by providing practical guides, training workshops, and repositories with examples tailored to various fields of knowledge. In line with this orientation, the universities emphasise the need to promote AI literacy and critical use of these tools, avoiding technological dependence and ensuring that AI acts as an ally in intellectual development rather than a shortcut that undermines deep learning (Gonsalves, 2024; Yan et al., 2024; Yusuf et al., 2024).

This approach is consistent with the principles outlined in the 'Blueprint for an AI Bill of Rights' (United States, 2022), which emphasizes the importance of individual autonomy, explainability, and accountability in relation to automated systems. Some universities, such as Oxford and UC Berkeley, differentiate expected uses according to academic level (undergraduate or postgraduate), promoting a more situated treatment of authorship and learning. Once again, Caltech appears with a score of 2, characterized by more restrictive policies and a narrower focus on transparency recommendations, but without a broadly structured institutional program.

In the information security axis, although all institutions presented at least minimal guidelines, eight obtained the maximum score (3), including Oxford, Harvard, Stanford, Cambridge, Princeton, UC Berkeley, Imperial College London, and Yale. These universities

explicitly align themselves with regulatory frameworks, such as the European Union's General Data Protection Regulation (GDPR) (European Union, 2016), making direct reference to the careful selection of platforms, restrictions on the use of open tools, and the protection of sensitive data related to staff, students, and institutional research. In these cases, one can observe the institutionalisation of academic cybersecurity practices, with the involvement of technical and legal sectors in defining protocols. At institutions such as MIT and Caltech, although important recommendations exist, they are presented in a less systematised way, reflecting an intermediate stage of normative development, with a primary focus on risk analysis and the use of licensed platforms.

The governance and ethics dimension presents the most heterogeneous picture in terms of institutional arrangements, despite all ten universities achieving the maximum score (3) on this axis. Common to these institutions are specific governance committees or equivalent structures, policies articulated in collaboration with research centres, and detailed guidelines on algorithmic justice, bias, auditability, and digital inclusion. At the same time, institutional designs vary: some universities concentrate governance in central bodies linked to the rectorate or provost, whereas others distribute responsibilities across multiple units, such as AI ethics centres, public policy institutes, and information technology services. In all cases, governance goes beyond mere compliance. It treats AI as an object of interdisciplinary reflection, in line with the UNESCO 'Recommendation on the Ethics of Artificial Intelligence' (2022), which calls for the involvement of multiple stakeholders in developing norms, with a focus on equity, human rights, and digital sustainability.

One interpretative point deserves emphasis. The universal presence of governance and ethics guidance should not be read as evidence of institutional uniformity. Rather, it suggests that universities already recognise generative AI as a matter of institutional legitimacy, not only pedagogical convenience. In this sense, ethical framing functions as a public demonstration that adoption will be mediated by principles of accountability, fairness, and human oversight (Porsdam Mann et al., 2024; UNESCO, 2022). What varies less is the acknowledgement of the issue than the organisational form through which it is addressed.

The comparison also reveals a tension between centralisation and distributed governance. More centralised arrangements tend to offer clearer institution-wide signals about disclosure, security, and approved uses. In contrast, distributed arrangements preserve disciplinary autonomy and allow schools, departments, and instructors to adapt guidance to local epistemic norms. Given this context, it is reasonable to interpret that the most promising models are hybrid: central institutions establish minimum standards for integrity, data protection, and transparency, while local units retain the flexibility to calibrate pedagogical use to disciplinary needs (Huttenlocher et al., 2023; Russell Group, 2023).

A final interpretative point concerns documentary visibility. Because this study relies on publicly available documents, the analysis captures the degree to which governance is formalised and communicated, not the full extent of internal practices. Universities with stronger communication ecosystems may therefore appear more robust, partly because their policies are more visible. This fact does not invalidate the comparative exercise; on the contrary, public codification is itself a relevant dimension of governance in contexts where students and staff need clear institutional signals. Still, the findings should be interpreted as a measure of the robustness of publicly documented guidance rather than a direct measure of implementation depth.

The integration of pedagogical policies, student guidelines, security strategies, and governance structures points to the emergence of an institutional model of response to generative AI that is both normative and formative. This model aligns with the principles

of accountability, transparency, and human-centered AI present in regulatory frameworks, such as the European Commission's proposal for AI regulation (European Commission, 2021), and it reinforces the role of universities as regulators of ethical practices in ecosystems increasingly mediated by algorithms.

By systematising these guidelines and quantifying their robustness through an analytical matrix, this study contributes to the delineation of an emerging field of applied research: the university governance of generative AI. The identification of patterns, gaps, and best practices among the analysed universities provides a comparative reference point for institutions in other contexts that seek to develop their own policies. The findings indicate that effective AI regulation in higher education requires more than disciplinary control mechanisms. It demands the strategic integration of technology, pedagogy, data protection, and ethical responsibility, guided by global frameworks and sensitive to local contexts.

Despite the analytical and applied contributions of this study, some limitations must be acknowledged. First, the institutional focus, restricted to the ten institutions with the best performance in THE 2025, concentrates on contexts with high technological density and advanced normative capacity, which may limit the generalisability of the findings to university settings in developing countries or in systems with lower levels of digital maturity. Second, documentary analysis, although methodologically rigorous, is conditioned by the public availability and transparency of institutional websites and may not capture practices under development or internal documents that have not been disclosed. Furthermore, the qualitative nature of coding, even when supported by explicit criteria and inter-rater validation, entails a degree of interpretation that could be enriched in future work through triangulation with interviews, observations, or in-depth case studies.

These limitations open space for promising research agendas. Comparative studies that incorporate universities from different geographical regions and institutional profiles, such as technical institutions, community colleges, or distance education universities, can broaden understanding of the diversity of responses to generative AI. Longitudinal research is also relevant for capturing normative evolution in the face of rapid technological change and regulatory cycles at national and international levels. Finally, investigations that combine the analysis of institutional guidelines with the perspectives of managers, faculty, and students may deepen our understanding of the implementation dynamics and the everyday challenges of algorithmic mediation in higher education.

IMPLICATIONS FOR INSTITUTIONAL POLICY AND PRACTICE

The comparative evidence allows the formulation of four practical recommendations for higher education institutions that are still consolidating their responses to generative AI.

First, institutions should combine centrally issued baseline policies with room for contextual adaptation at programme, department, or course level. Central guidance is especially important for disclosure, data protection, and minimum integrity standards, whereas local adaptation is crucial for discipline-specific pedagogical uses.

Second, universities should formalise disclosure rules for students and staff, specifying when acknowledgement is mandatory, what exactly must be disclosed, and how AI-assisted work should be evaluated. Clear transparency rules reduce ambiguity and make accountability more enforceable.

Third, governance frameworks should clearly distinguish between approved and non-approved tools and explicitly prohibit the insertion of sensitive, personal, or research-restricted data into public platforms. This context depicts one area where institutional clarity has the greatest immediate risk-reduction value.

Fourth, universities should invest in continuous AI literacy programmes for students, faculty, and professional staff. The comparative evidence suggests that robust governance is not achieved through prohibition alone, but through training, guidance, and ethically informed use.

CONCLUSION

In a global context marked by the accelerating pace of technological innovation, generative AI poses new challenges and opportunities for higher education institutions. This study sought to understand how the ten higher education institutions with the best performance in the Times Higher Education (THE) World University Rankings 2025 (Times Higher Education, 2025) have structured their institutional guidelines on the use of generative AI, based on documentary analysis of 87 official documents systematised into four thematic axes: teaching, student use, information security and governance/ethics.

The results indicate that these universities operate not only as regulatory spaces but also as formative instances of a new academic ethos oriented towards critical reflection, transparency, and shared responsibility in the use of AI. The presence of robust guidelines articulated across multiple institutional sectors reveals an integrated governance model that recognises the transversal nature of AI and its capacity to reconfigure pedagogical, normative, and administrative practices. Rather than adopting technologies in a merely instrumental way, these institutions lead a process of normative and formative institutionalisation of generative AI, grounded in values such as equity, algorithmic justice, data protection, and digital citizenship. The findings demonstrate broad adherence to the formulation of consistent AI governance policies, indicating an emerging pattern of good practices that can guide institutions in diverse contexts.

By systematising and critically analysing such guidelines, the study offers theoretical and practical contributions to institutions at the initial stages of policy development on AI, helping to build comparative frameworks and internationally recognisable good practices. The consolidation of institutional frameworks aligned with global ethical principles, such as those advocated by UNESCO, the European Union, and the United States' AI Bill of Rights, is essential to ensure that generative AI is incorporated into the university mission without undermining its formative, critical, and humanistic dimensions.

Finally, it is reaffirmed that the future of higher education in AI-mediated contexts will depend less on unrestricted access to technology and more on institutional capacity to mobilise it in ethical, transparent, and public-good-oriented ways. In this new historical cycle, the university is called upon not only to act as a user, but as a protagonist in the formulation of norms, values, and practices that will shape educational horizons in an algorithmically expanded society.

AUTHORS CONTRIBUTIONS

Conceptualization: Camila Lopes Ferreira; Data curation: Camila Lopes Ferreira; Formal analysis: Camila Lopes Ferreira; Investigation: Camila Lopes Ferreira; Methodology: Camila Lopes Ferreira; Project administration: José Roberto Herrera Cantorani; Supervision: Claudia Tania Picinin; Validation: Alessandra Dutra and Sani de Carvalho Rutz da Silva; Writing – original draft: Camila Lopes Ferreira; Writing – review & editing: Alessandra Dutra, Sani de Carvalho Rutz da Silva, José Roberto Herrera Cantorani and Claudia Tania Picinin.

REFERENCES

AI at Yale. (2025a). *Clarity Platform*. https://ai.yale.edu/yales-ai-tools-and-resources/clarity-platform

AI at Yale. (2025b). *Envisioning AI at Yale: An interdisciplinary symposium*. https://ai.yale.edu/symposium-2025

Berkman Klein Center for Internet & Society. (2025). *Ethics and governance of AI initiative*. https://cyber.harvard.edu/projects/ethics-and-governance-ai-initiative

Bowen, G. A. (2009). Document analysis as a qualitative research method. *Qualitative Research Journal, 9*(2), 27-40. https://doi.org/10.3316/QRJ0902027

California Institute of Technology. (2025a). *Resources for teaching in the age of AI*. https://ctlo.caltech.edu/universityteaching/resources/resources-for-teaching-in-the-age-of-ai

California Institute of Technology. (2025b). *HSS policy on generative AI*. https://www.hss.caltech.edu/hss-policies/hss-policy-on-generative-ai

California Institute of Technology. (2025c). *Guidelines for secure and ethical use of artificial intelligence (AI)*. https://www.imss.caltech.edu/services/ai/caltech-ai/guidelines-for-secure-and-ethical-use-of-artificial-intelligence-ai

California Institute of Technology. (2025d). *The Linde Center for Science, Society, and Policy (LCSSP)*. https://lindeinstitute.caltech.edu/research/lcssp

California Institute of Technology. (2025e). *MACH 33: The Caltech Festival of New Science-Driven Plays*. https://tacit.caltech.edu/mach33

California Institute of Technology. (2025f). *Undergraduate and graduate admissions AI policy: Guidelines for fall applicants*. https://www.admissions.caltech.edu/apply/first-year-applicants/supplemental-application-essays/ethical-use-of-ai-guidelines-for-fall-applicants

Chen, X., & Wu, D. (2024). Automatic generation of multimedia teaching materials based on generative AI: Taking Tang poetry as an example. *IEEE Transactions on Learning Technologies, 17*, 1327-1340. https://doi.org/10.1109/TLT.2024.3378279

Cowen, R. (2013). Changing principles and goals of universities: Questioning trajectories. *Sisyphus - Journal of Education, 1*(2), 38-53. https://doi.org/10.25749/sis.3627

Cress, U., & Kimmerle, J. (2023). Co-constructing knowledge with generative AI tools: Reflections from a CSCL perspective. *International Journal of Computer-Supported Collaborative Learning*, 18, 607-614. https://doi.org/10.1007/s11412-023-09409-w

European Commission. (2021). *Proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts*. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206

European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union, L 119*, 1-88. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679

Gonsalves, C. (2024). Generative AI's impact on critical thinking: Revisiting Bloom's taxonomy. *Journal of Marketing Education*. Advance online publication. https://doi.org/10.1177/02734753241305980

Harvard Medical School. (2025). *Guidance and tools for secure, ethical use of AI at HMS*. https://it.hms.harvard.edu/about/policies-and-guidelines/generative-ai

Harvard University Information Technology. (2025). *Generative artificial intelligence (AI)*. https://www.huit.harvard.edu/ai

Harvard University Office of the Provost. (2025). *Guidelines for using ChatGPT and other generative AI tools at Harvard*. https://provost.harvard.edu/guidelines-using-chatgpt-and-other-generative-ai-tools-harvard

Harvard University. (2025). *AI ethics in business: Managing bias and ethical usage*. https://pll.harvard.edu/course/ai-ethics-business-managing-bias-and-ethical-usage

Huttenlocher, D., Ozdaglar, A., & Goldston, D. (2023). *A framework for U.S. AI governance: Creating a safe and thriving AI sector*. MIT Schwarzman College of Computing; MIT Washington Office. https://computing.mit.edu/wp-content/uploads/2023/11/AIPolicyBrief.pdf

Imperial College London. (2025a). *Generative AI tools guidance*. Library Services. https://www.imperial.ac.uk/admin-services/library/learning-support/generative-ai-guidance/

Imperial College London. (2025b). *AI for teaching*. Teaching Toolkit. https://www.imperial.ac.uk/staff/educational-development/teaching-toolkit/use-of-generative-ai-for-teaching-learning-and-assessment/ai-for-teaching/

Imperial College London. (2025c). *AI fundamentals*. Imperial Policy Forum. https://www.imperial.ac.uk/research-and-innovation/the-forum/courses-for-policymakers/fundamentals-/ai-fundamentals/

Kong, S.-C., & Yang, Y. (2024). A human-centered learning and teaching framework using generative artificial intelligence for self-regulated learning development through domain knowledge learning in K-12 settings. *IEEE Transactions on Learning Technologies*, 17, 1562-1573. https://doi.org/10.1109/TLT.2024.3392830

Massachusetts (2025a). *201 CMR 17.00*: Standards for the protection of personal information of residents of the Commonwealth. https://www.mass.gov/regulations/201-CMR-1700-standards-for-the-protection-of-personal-information-of-residents-of-the-commonwealth

Massachusetts (2025b). *General Laws, Chapter 93H*: Security breaches. https://malegislature.gov/Laws/GeneralLaws/PartI/TitleXV/Chapter93H

Mishra, P., Oster, N., & Henriksen, D. (2024). Generative AI, teacher knowledge and educational research: Bridging short- and long-term perspectives. *TechTrends*, 68, 205-210. https://doi.org/10.1007/s11528-024-00938-1

MIT Information Systems and Technology. (2025). *Guidance for use of generative AI tools*. https://ist.mit.edu/ai-guidance

MIT Internet Policy Research Initiative. (2024). *AI policy: 2024*. https://internetpolicy.mit.edu/ai-policy/

MIT Management. (2025). *Getting started with AI-enhanced teaching: A practical guide for instructors*. https://mitsloanedtech.mit.edu/ai/teach/getting-started/

MIT School of Architecture + Planning. (2023). *Ethics and governance of artificial intelligence*. https://www.media.mit.edu/groups/ethics-and-governance/overview/

Mittal, U., Sai, S., Chamola, V., & Sangwan, D. (2024). A comprehensive review on generative AI for education. *IEEE Access*, 12, 142733-142759. https://doi.org/10.1109/ACCESS.2024.3468368

Payne, B. H. (2019). *An ethics of artificial intelligence curriculum for middle school students*. https://thecenter.mit.edu/wp-content/uploads/2020/07/MIT-AI-Ethics-Education-Curriculum.pdf

Porsdam Mann, S., Vazirani, A. A., Aboy, M., Earp, B. D., Minssen, T., Cohen, I. G., & Savulescu, J. (2024). Guidelines for ethical use and acknowledgement of large language models in academic writing. *Nature Machine Intelligence*, 6, 1272-1274. https://doi.org/10.1038/s42256-024-00922-7

Princeton University. (2025a). *Princeton dialogues on AI and ethics - Case study PDFs*. https://aiethics.princeton.edu/case-studies/case-study-pdfs/

Princeton University. (2025b). *Generative AI and Our Classrooms*. McGraw Center for Teaching and Learning. https://mcgraw.princeton.edu/generative-ai

Princeton University. (2025c). *Generative AI at Princeton*. Office of Information Technology. https://oit.princeton.edu/generative-ai

Princeton University. (2025d). *Microsoft Copilot at Princeton University*. Office of Information Technology. https://oit.princeton.edu/microsoft-copilot-princeton-university

Princeton University. (2025e). *Princeton-supported AI tools*. Office of Information Technology. https://oit.princeton.edu/generative-ai/ai-tools

Princeton University. (2025f). *Princeton dialogues on AI and ethics*. https://aiethics.princeton.edu/about/

Princeton University Library. (2025). *Disclosing the use of generative AI at Princeton University*. https://libguides.princeton.edu/generativeAI/disclosure

Ramirez, F. O. (2013). World society and the university as formal organization. *Sisyphus - Journal of Education, 1*(1), 124-153. https://doi.org/10.25749/sis.2833

Russell Group. (2023). *Principles on the use of generative AI tools in education*. https://www.russellgroup.ac.uk/policy/policy-briefings/principles-use-generative-ai-tools-education

Stanford HAI. (2025). *The 2025 AI Index Report*. Stanford Institute for Human-Centered AI. https://hai.stanford.edu/ai-index/2025-ai-index-report

Stanford University. (2025a). *AI meets education at Stanford*. Stanford Center for Teaching and Learning. https://ctl.stanford.edu/aimes

Stanford University. (2025b). Report outlines Stanford principles for use of AI. *Stanford Report*. https://news.stanford.edu/stories/2025/01/report-outlines-stanford-principles-for-use-of-ai

Stanford University. (2025c). *Artificial intelligence teaching guide*. Stanford Teaching Commons. https://teachingcommons.stanford.edu/teaching-guides/artificial-intelligence-teaching-guide

Stanford University. (2025d). *Course policies on generative AI use*. Teaching and Learning Hub. https://tlhub.stanford.edu/docs/course-policies-on-generative-ai-use/

Stanford University. (2025e). *Artificial Intelligence (AI)*. https://teachingcommons.stanford.edu/news/artificial-intelligence-ai

Times Higher Education. (2025). *World University Rankings 2025*. https://www.timeshighereducation.com/world-university-rankings/2025/world-ranking

UNESCO. (2022). *Recommendation on the ethics of artificial intelligence*. United Nations Educational, Scientific and Cultural Organization. https://unesdoc.unesco.org/ark:/48223/pf0000381137

United States, Office of Science and Technology Policy. (2022). *Blueprint for an AI Bill of Rights: Making automated systems work for the American people*. The White House. https://www.govinfo.gov/app/details/GOVPUB-PREX23-PURL-gpo193638

University of California, Berkeley. (2025a). *Guidance on the use of AI*. Berkeley Ai Hub. https://ai.berkeley.edu/guidance-use-ai

University of California, Berkeley. (2025b). *AI Policy Hub*. Center for Long-Term Cybersecurity. https://cltc.berkeley.edu/program/ai-policy-hub/

University of California, Berkeley. (2025c). *Appropriate use of generative AI tools*. Office of Ethics, Risk & Compliance Services. https://oercs.berkeley.edu/appropriate-use-generative-ai-tools

University of California, Berkeley. (2025d). *AI in teaching & learning overview*. Research, Teaching, and Learning. https://rtl.berkeley.edu/ai-teaching-learning-overview

University of Cambridge. (2025a). *Generative AI and assessment*. Blended Learning Service. https://blendedlearning.cam.ac.uk/artificial-intelligence-and-education/generative-ai-and-assessment

University of Cambridge. (2025b). *Using generative AI*. Blended Learning Service. https://blendedlearning.cam.ac.uk/artificial-intelligence-and-education/using-generative-ai

University of Cambridge. (2025c). *How we use generative AI tools*. Communications. https://www.communications.cam.ac.uk/generative-ai-tool-guidelines

University of Cambridge. (2025d). *AI: Futures and responsibility*. Leverhulme Centre for the Future of Intelligence. https://www.lcfi.ac.uk/research/programme/ai-futures-and-responsibility

University of Oxford. (2025a). *AI and academic practice*. Centre for Teaching and Learning. https://www.ctl.ox.ac.uk/ai

University of Oxford. (2025b). *Generative AI guidance for communicators*. Communications Resources. https://communications.admin.ox.ac.uk/ai-guidance

University of Oxford. (2025c). *Institute for Ethics in AI*. https://www.oxford-aiethics.ox.ac.uk/

University of Oxford. (2025d). *Use generative AI services safely*. https://www.infosec.ox.ac.uk/use-generative-ai-services-such-as-chatgpt-safely

University of Oxford. (2025e). *Guidance on safe and responsible use of GenAI*. https://www.ox.ac.uk/students/life/it/genai-tools/guidance-on-safe-and-responsible-use-of-genai

U.S. Department of Education (2011). *Family Educational Rights and Privacy Act* (*FERPA)*. https://studentprivacy.ed.gov/ferpa

U.S. Department of Health & Human Services (2026). *Health Information Privacy*. https://www.hhs.gov/hipaa/index.html

Yale University. (2023). *Guidelines for the use of generative AI tools*. Office of the Provost. https://provost.yale.edu/news/guidelines-use-generative-ai-tools

Yale. (2025a). *AI guidance for teachers*. Poorvu Center for Teaching and Learning. https://poorvucenter.yale.edu/teaching/teaching-resource-library/ai-guidance-for-teachers

Yale. (2025b). *AI Guidance for Students*. Poorvu Center for Teaching and Learning. https://poorvucenter.yale.edu/learning/learning-resource-library/ai-guidance-for-students

Yan, L., Greiff, S., Teuber, Z., & Gašević, D. (2024). Promises and challenges of generative artificial intelligence for human learning. *Nature Human Behaviour*, 8, 1839-1850. https://doi.org/10.1038/s41562-024-02004-5

Yusuf, A., Pervin, N., Román-González, M., & Noor, N. M. (2024). Generative AI in education and research: A systematic mapping review. *Review of Education, 12*(2), e3489. https://doi.org/10.1002/rev3.3489

Zhai, X. (2024). Transforming teachers' roles and agencies in the era of generative AI: Perceptions, acceptance, knowledge, and practices. *Journal of Science Education and Technology*. Advance online publication. https://doi.org/10.1007/s10956-024-10174-0

*

Received: November 21, 2025
Revisions Required: March 3, 2026
Accepted: April 8, 2026
Published online: June 30, 2026