HEA | An tÚdarás um Ard-Oideachas
The Higher Education Authority

Generative AI in Higher Education Teaching & Learning

Policy Framework

Contributors
James O'Sullivan
Colin Lowry
Ross Woods
Tim Conlon

Acknowledgements
Alan Smeaton (Government of Ireland AI Advisory Council), Barry O'Sullivan (Government of Ireland AI Advisory Council), Susan Leavy (Government of Ireland AI Advisory Council), Caoimhe Hope (Department of Further and Higher Education, Research, Innovation and Science), Anne Ribault-O'Reilly (Department of Further and Higher Education, Research, Innovation and Science), Ana Rocha (Higher Education Authority), Imma Zoppi (Higher Education Authority), Paul O'Donovan (University College Cork), Joseph Feller (University College Cork), Danielle Duignan (AHEAD), Jim O'Mahony (Munster Technological University), Denise Mac Giolla Rí (Technological University of the Shannon), Justin Tonra (University of Galway), Brian Marrinan (Journey Partners), Clyde Hutchinson (Journey Partners), Mary-Claire Kennedy (University of Limerick), Maria Murphy (Maynooth University), Derek Dodd (Technological University Dublin), Pauline Rooney (Trinity College Dublin), Loretta Goff (University College Cork), Roisín Morris-Drennan (Quality and Qualifications Ireland), Liam Fogarty (University College Dublin), Barbara Whelan (Microsoft), Johanna Archbold (Atlantic Technological University), Áine Clarke (Ibec), Marie Clarke (University College Dublin), Bryan O'Mahony (Aontas na Mac Léinn in Éirinn), Nessa McEniff (Learnovate), Rosemary Day (Mary Immaculate College), Leo Casey (National College of Ireland), Alison Cook-Sather (Bryn Mawr College), Jan McArthur (Lancaster University), Paul McSweeney (University College Cork), Áine Ní Shé (Munster Technological University), Frances O'Connell (Technological University of the Shannon), Emma Muldoon Ryan (Aontas na Mac Léinn in Éirinn), Tim Thompson (Maynooth University), Garrett Murray (Enterprise Ireland), Niamh Kennedy (National Student Engagement Programme), Buster Whelan (Trinity College Dublin Students' Union), Eamon Costello (Dublin City University), Ann Riordan (University College Cork).

HEA Generative AI Policy Framework
https://hub.teachingandlearning.ie/genai/policy-framework

HEA Generative AI Resource Portal
https://hub.teachingandlearning.ie/genai/

Generative AI in Higher Education Teaching & Learning: Policy Framework

Version 1.0, December 2025
DOI: 10.82110/073e-hg66

Higher Education Authority, Dublin

How to cite:
O'Sullivan, James, Colin Lowry, Ross Woods & Tim Conlon. *Generative AI in Higher Education Teaching & Learning: Policy Framework.* Higher Education Authority, 2025. DOI: 10.82110/073e-hg66.

Table of Contents

Purpose of the Framework

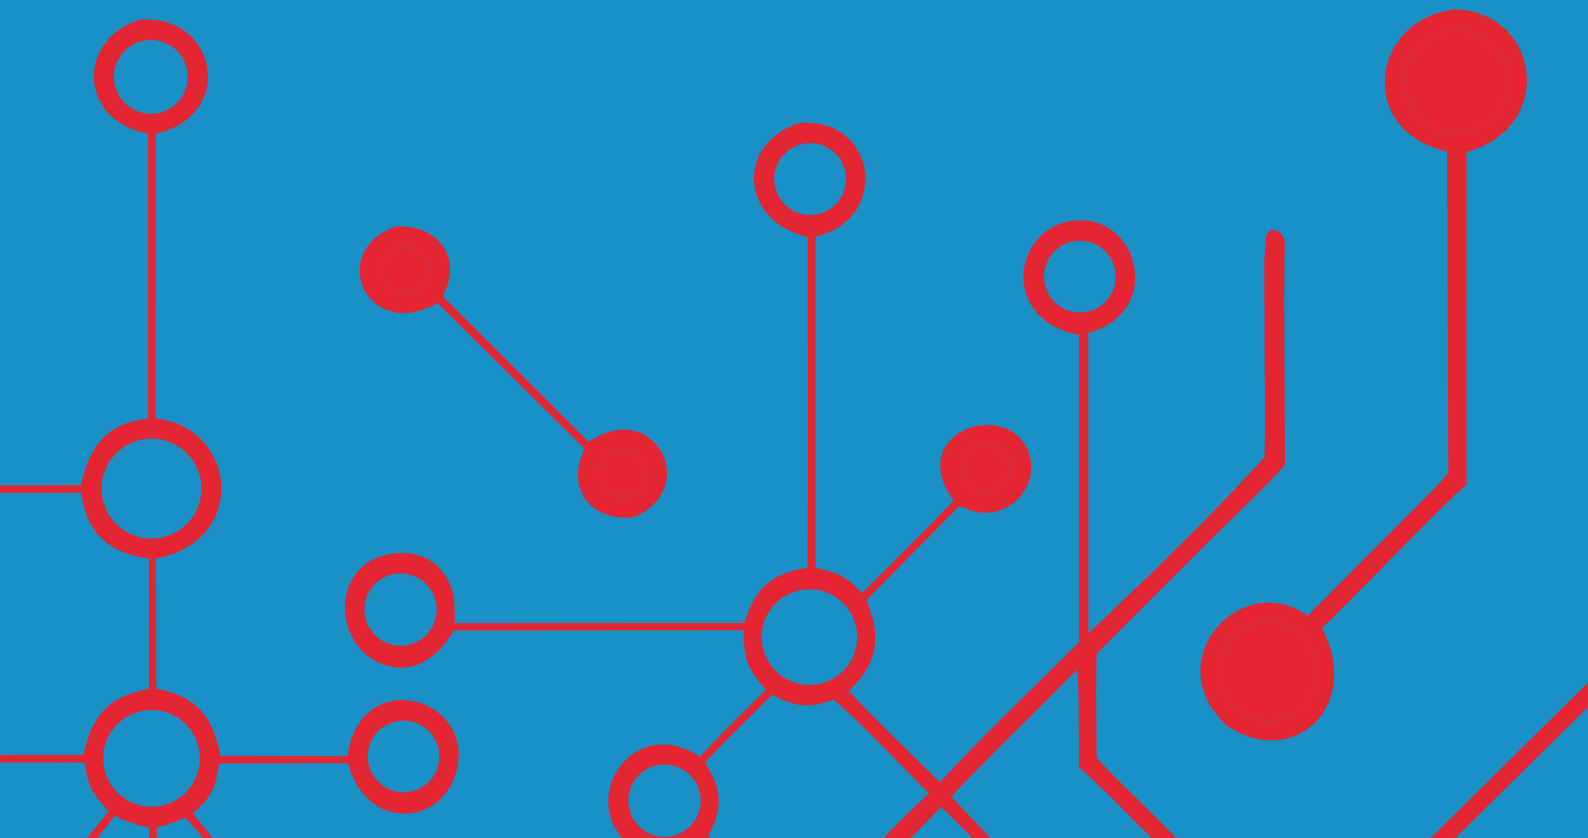

This policy framework, along with its and supporting instruments,[1] applies to the use of generative artificial intelligence (gen AI), most notably large language models (LLMs) like ChatGPT, in teaching and learning within Irish higher education institutions (HEIs). Its purpose is to guide educators, academic leaders, and professional staff in making informed, values-based decisions about how gen AI is adopted and integrated into educational practice.

This framework deals specifically with generative AI, recognising that while gen AI forms part of the broader field of artificial intelligence, it raises distinctive opportunities and risks for higher education. It is therefore important to distinguish between 'AI' in general and 'gen AI' as the focus of this policy. Generative artificial intelligence refers to systems that can produce new content, such as text, images, or code, in response to prompts, based on patterns learned from large datasets. Large language models like ChatGPT, Copilot, and Claude are the most prominent example.

The focus of this policy framework is specifically on teaching and learning. While gen AI has implications across research, administration, and institutional strategy, this framework is explicitly concerned with how such technologies reshape learning design, pedagogy, student engagement, assessment, and academic integrity. It is intended as a tool for reflection and structured decision-making within these domains. The HEA's decision to focus this initial framework on teaching and learning reflects the immediacy and scale of the impact that generative AI is already having on students and educators. The pedagogical sphere has become the most visible site of disruption, where questions of academic integrity, assessment design, and equitable participation have demanded urgent attention. This focus does not seek to reinforce an artificial separation between research and teaching, which are fundamentally interdependent within higher education. Rather, it recognises that the classroom is where the implications of generative AI are most acutely felt, and that insights developed here will inevitably inform broader institutional and research practices.

[1] See https://hub.teachingandlearning.ie/genai/policy-framework

There is no intention to prescribe a uniform set of rules or single model of adoption. Instead, this policy framework sets out a values-based orientation that can guide institutions in developing their own policies, practices, and cultures of use. Frameworks of this kind provide coherence at system level, establish principles that can be adapted and operationalised locally, and support national coordination without constraining institutional autonomy or innovation.

This framework does not apply directly to students, nor is it intended as guidance for their individual use of AI tools. It provides direction to those responsible for designing, delivering, and supporting teaching and learning. By doing so, it aims to shape the conditions in which students encounter AI in their education, ensuring that institutional practices are coherent, ethical, and pedagogically sound.

Within these defined boundaries, this framework provides a national reference point that can be adapted to the specific contexts of individual HEIs, supporting coherent system-wide engagement with gen AI while respecting institutional autonomy.

In this context, this policy framework seeks to:

(1) Provide HEIs with a structured but adaptable set of values to underpin institutional decision-making on gen AI;

(2) Encourage responsible and pedagogically meaningful adoption of gen AI that safeguards the interests of students and staff;

(3) Promote national coherence while enabling institutional autonomy and innovation; and

(4) Position Irish higher education as a leader in the responsible and values-driven adoption of gen AI.

This framework takes as its starting point the position that gen AI is neither a passing novelty nor a universal remedy. It is a set of tools that, regardless of any individual professional or personal perspective, must be integrated thoughtfully into teaching and learning in ways that are consistent with academic values, national policy commitments, and the lived realities of HEIs in Ireland.

This framework reflects the current state of generative AI adoption in higher education teaching and learning as of the date of publication, and is intended to evolve in response to technological developments, emerging evidence, and sector experience. The HEA will issue updates as and when required, and institutions should refer to the most recent published version when developing or reviewing their own policies.

Context

The rapid evolution of generative artificial intelligence is reshaping higher education worldwide.[2] These technologies offer significant potential to enhance student learning, enabling more responsive forms of teaching and strengthening institutional operations. They also present recognised risks and uncertainties in bias and inaccuracy, data protection and intellectual property concerns, environmental impact, inequitable access, threatening fundamental academic values such as transparency and fairness.

National and international guidance on generative AI in education, frameworks emerging from peer higher education systems, and research on the pedagogical implications of AI all contribute to the evidence base on which this policy framework has been developed. Specific attention has been given to ensuring alignment with existing work by the Government of Ireland's AI Advisory Council, established to provide strategic guidance on AI adoption across all sectors, including education;[3] the work of Quality and Qualifications Ireland (QQI) on academic integrity and quality assurance in an AI-enabled environment; the European Union's AI Act and evolving regulatory framework, as well as the guidance from the European Commission's High-Level Expert Group on Artificial Intelligence on ethical AI development and deployment;[4] and international frameworks from bodies such as UNESCO and the OECD on AI in education, competency development, and educational equity.[5]

This policy framework also draws on a growing body of international, peer-reviewed scholarship examining the pedagogical, ethical, and institutional implications of generative AI in higher education.[6] This research provides critical insights into how AI is reshaping higher education teaching and learning practices and offers comparative perspectives that have guided the development of this framework.

There is a substantial body of national sectoral evidence, including a survey of staff and students conducted by QQI[7] and a major national consultation undertaken by the Higher Education Authority (HEA).[8] These sources provide a nuanced picture of how gen AI is being encountered across higher education in Ireland. They show that adoption is uneven but rapidly growing, and that many students already use AI routinely for content generation, while staff are beginning to explore applications in teaching support, feedback, and formative assessment. However, despite rapid adoption, there are significant concerns about academic integrity, the reliability of detection tools, and the potential for inequities in access and use.

2 Wang et al., 'Generative AI in Higher Education: Seeing ChatGPT through Universities' Policies, Resources, and Guidelines'; Deng et al., 'Does ChatGPT Enhance Student Learning? A Systematic Review and Meta-Analysis of Experimental Studies.'

3 Smeaton et al., 'AI and Education.'

4 'Ethics Guidelines for Trustworthy AI.'

5 Miao et al., 'AI and Education: Guidance for Policy-Makers'; 'Ethical Guidelines on the Use of Artificial Intelligence (AI) and Data in Teaching and Learning for Educators'; Lodge et al., 'Assessment Reform for the Age of Artificial Intelligence'; Miao and Holmes, 'Guidance for Generative AI in Education and Research'; Varsik and Vosberg, 'The Potential Impact of Artificial Intelligence on Equity and Inclusion in Education'; Miao and Mutlu, 'AI Competency Framework for Teachers'; Elhussein et al., 'Shaping the Future of Learning: The Role of AI in Education 4.0'; Miao et al., 'AI Competency Framework for Students'; Hoernig et al., 'Generative AI and Higher Education'; Hemment and Kommers, 'Doing AI Differently.'

6 For further information on the research and policy evidence underpinning this framework, see supporting instruments.

7 *Analysis of Results from the Generative Artificial Intelligence Survey 2025.*

8 O'Sullivan et al., *Generative AI in Higher Education Teaching and Learning: Sectoral Perspectives.*

The HEA's national focus group process identified that institutional responses to AI have, to date, tended to be fragmented and largely driven by individual staff initiative rather than coordinated strategic approaches. Assessment has emerged as the most immediate area of pressure, with long-standing questions about its purpose and design brought into sharper focus by the advent of AI. Participants also highlighted the risk that issues of equity and inclusion could be exacerbated in the absence of clear, sector-wide guidance. Importantly, many focus groups emphasised that generative AI has not created these challenges, but has intensified pre-existing concerns within higher education.

Staff highlighted significant workload pressures, noting that concerns about the additional demands of adapting to AI are layered onto already substantial responsibilities for teaching, feedback, and student support. Students described persistent challenges in maintaining engagement, with AI further complicating how participation and effort are recognised and rewarded.

Long-standing discussions about the skills and capabilities developed through higher education, beyond those articulated in formal learning outcomes, have been brought into sharper focus by the ability of AI systems to replicate many routine academic processes. Similarly, discourse on assessment design and fairness, which have historically been articulated in relation to plagiarism, group work, and grading practices, are now being reframed by the emergence of AI systems capable of producing outputs that mirror authentic student submissions.

The framework is aligned with international standards for trustworthy and ethical AI, and with broader Irish public service commitments to responsible AI adoption. It is designed to promote coherence at the system level while enabling HEIs to develop strategies that reflect their own institutional priorities, cultures, and communities.

While many international guidelines on ethical AI are drafted primarily with system designers and developers in view, their relevance extends well beyond the point of technical production. They offer higher education institutions a coherent vocabulary and evaluative structure for thinking through the ethical and pedagogical implications of adopting particular AI systems. Read in this way, existing guidelines on trustworthy and ethical AI enable educators and institutional leaders to scrutinise not only what a tool can do, but how its designers have interpreted notions such as robustness, transparency, accountability, and human oversight. They also make visible the trade-offs that vendors may have accepted in pursuit of performance, scale, or market advantage, and invite reflection on how those compromises might surface in teaching, assessment, and academic governance. Even where guidelines have been produced specifically for developers, familiarity with these frameworks supports a more critical form of adoption, one grounded less in novelty or efficiency gains than in an informed judgement about whether a system's conception of trustworthiness aligns with the values and obligations of educational practice.

Whether viewed with enthusiasm or caution, higher education now operates in a world where generative artificial intelligence is embedded in the everyday tools of students and staff. Tools capable of producing fluent written text and other outputs are no longer peripheral but increasingly woven into browsers and productivity tools. Traditional assumptions about student work and assessment can no longer be sustained without significant reconsideration.

Debates about gen AI in higher education are often marked by polarised positions. Some staff regard gen AI as a fundamental threat to academic integrity and critical and creative development that should be tightly controlled or excluded, while others present it as a ready solution to enduring challenges such as workload pressures and student engagement to be adopted without reservation. In practice, neither extreme provides a sufficient basis for guiding sector-wide action.

If higher education is to continue to privilege human judgement as central to scholarship and professional practice, there must be scope to create learning contexts where reasoning, creativity, and critical voice are demonstrated independently of generative AI. In such cases, the decision to limit or prohibit the use of AI should be understood as a deliberate pedagogical choice, undertaken within a framework of critical engagement and transparency so that students understand its rationale and purpose.

For higher education, the challenge now is to reaffirm and uphold core educational purposes, such as the cultivation of critical thinking, independent reasoning, and intellectual integrity, while recognising that some long-standing practices are now under pressure from emerging technologies. In particular, reliance on conventional assessment formats that treat written work as the primary measure of learning may need reconsideration in light of technologies that can replicate such outputs. This does not diminish the value of assessment, nor the necessity of awarding degrees, credits, and grades, but rather, calls for a shift away from transactional models of evidence towards approaches that more authentically reflect the values and capabilities higher education seeks to foster in students. Achieving this will be a complex and iterative process.

The HEA advocates a values-based approach to gen AI adoption, which requires moving beyond both uncritical adoption and uncritical rejection. Categorical prohibition risks ignoring the inevitability of student and staff engagement with these tools, while unqualified enthusiasm risks overlooking real challenges, including bias, equity of access, environmental costs, and the safeguarding of critical judgement. Effective leadership and practice therefore lie in the space between, in acknowledging risks while recognising opportunities, and developing pedagogical and policy responses that are proportionate, evidence-informed, and grounded in the moral and intellectual missions of higher education teaching and learning.

Framework Principles[9]

[9] Each of the principles is further expanded in the accompanying document, Generative AI in Higher Education Teaching & Learning: Principles for Ethical AI Adoption.

These five principles provide an ethical and pedagogical foundation for how higher education institutions approach generative AI. They translate practical considerations into enduring commitments that should guide institutional decisions about adoption, governance, and use. While gen AI technology will continue to evolve rapidly, these principles anchor AI integration in the values that define higher education as a public good: the pursuit of knowledge with integrity, the commitment to equity and inclusion, the preservation of human judgment and agency, the protection of privacy and institutional autonomy, and the responsibility to sustain both educational practice and the environment for future generations. They are not prescriptive rules but rather a shared framework that allows HEIs to navigate complexity with clarity of purpose, ensuring that innovation serves rather than compromises the mission of higher education in Ireland.

Principle 1: Academic Integrity, Transparency, and Accountability

The university's role as guardian of legitimate knowledge depends on transparent disclosure of how that knowledge is produced. When gen AI contributes to intellectual work, the provenance of ideas matters to their evaluation and development. Accountability requires that human agents remain traceable and responsible for what is certified, regardless of the tools used in its creation. The principle acknowledges that gen AI has fundamentally altered what 'original work' can mean, but insists that this transformation demands greater rather than lesser clarity about authorship, contribution, and verification. HEIs are called to create environments where honest engagement with gen AI is rewarded and where students understand that the intellectual labour of critique and judgment cannot be outsourced. The fundamental principle that students are responsible for any work they submit remains unchanged. The aim is not to eliminate gen AI from academic work but to ensure that its use is visible and consonant with the development of genuine understanding, while equally protecting the right of staff to decline its use where they judge it inappropriate to their learning or teaching. Academic integrity in an AI age means preserving the relationship between learning and effort while acknowledging that the nature of that effort has changed.

Principle 2: Equity & Inclusion

Gen AI adoption must not become another mechanism through which existing inequalities are reproduced or amplified. Access to the benefits of AI-enabled education should be equitable and must not be determined by factors such as private means, linguistic background, or prior technical confidence. HEIs bear responsibility for ensuring that the tools which shape learning opportunities are available to all students on fair terms, and that the systems themselves do not embed biases that disadvantage particular communities. The principle recognises that equity extends beyond access to encompass how gen AI systems are trained, how they perform across languages and cultures, and whether their use in assessment privileges certain styles of expression or modes of engagement. It calls HEIs to attend to intersectional disadvantage and to understand that technology, left ungoverned, tends to favour the already advantaged. The spirit of this principle is one of active remedy rather than passive neutrality, such that institutions should design AI adoption with inclusion at its centre, not as an afterthought, and should be prepared to decline tools or practices that cannot be made equitable.

Principle 3: Critical Engagement, Human Oversight, & AI Literacy

Human judgment and agency must remain at the centre of educational practice, even as gen AI systems become more capable and persuasive. It recognises that education is fundamentally about developing the capacity to think, question, and create, and that these capabilities atrophy when algorithmic outputs are accepted without scrutiny. AI literacy is therefore an intellectual and ethical stance, equipping students with the ability to interrogate how systems work, to recognise their limits and biases, and to situate their use within broader considerations of knowledge and power in educational and societal contexts. Human oversight ensures that decisions affecting students' lives and futures remain the responsibility of people who can be held accountable and who understand the moral weight of those decisions. The principle cautions against the assumption that efficiency or scale alone justify delegating human judgement to automated systems. Instead, it expects HEIs to cultivate discernment in both staff and students, ensuring that AI serves educational aims rather than displacing them, and that graduates emerge not as skilled users of black-box tools but as thoughtful actors capable of shaping technology's role in their professions and societies.

Principle 4: Privacy & Data Governance

Educational data is uniquely sensitive, revealing intimate patterns of intellectual development, challenge, and achievement that require the highest standards of protection. It recognises that students are often in vulnerable positions when they generate data through learning activities, and that institutions hold these records in trust rather than as property to be exploited. The spirit of the principle is one of stewardship, that institutions are custodians of data that must be handled with transparency and respect for the autonomy and dignity of the individuals it represents. Ideally, Ireland's higher education institutions and sector should retain sovereignty over their data and systems, maintaining meaningful control over educational records and decision-making, and ensuring that partnerships with external providers do not compromise institutional autonomy or public accountability. Data ownership should rest with the institution or the student, and institutions should retain the capacity to audit the systems they deploy and, where necessary, to withdraw from them without loss of access to records or compromise to their educational mission. The principle asserts that neither efficiency nor innovation can justify compromises to privacy, institutional governance, or students' autonomy over the intellectual data produced through their participation in higher education. It acknowledges that generative AI systems often rely on data-intensive processes that can introduce risks of exposure, misuse, or bias, and it encourages higher education institutions to establish governance arrangements that prioritise data protection and institutional oversight over convenience or vendor assurances.

Principle 5: Sustainable Pedagogy

AI adoption should enhance, rather than diminish, the long-term sustainability of educational practice and the environment in which it operates. Sustainability encompasses both ecological responsibility and educational vitality, and so higher education institutions should account for the carbon costs of computation while also ensuring that generative AI does not erode the intellectual capacities that make human learning meaningful. This principle reflects a commitment to stewardship across time, recognising that education is not solely about efficiency or performance in discrete tasks, but about nurturing enduring qualities of curiosity, resilience, and independent thought. Where gen AI displaces the productive struggle through which these qualities develop, it risks undermining the very purposes it seeks to support. The principle also acknowledges that technological dependence carries risks, including vendor lock-in, infrastructure fragility, and the homogenisation of pedagogical practice. Institutions are therefore encouraged to maintain diversity in how they teach and assess, to preserve opportunities where learning happens without algorithmic mediation, and to ensure that the integration of gen AI remains a considered choice rather than an unexamined default. Sustainable pedagogy requires that the use of generative AI be capable of being maintained, adapted, or, where necessary, discontinued without compromising the institution's capacity to deliver on its educational mission.

Operationalising the Principles

While a broader suite of recommendations is detailed in supporting documents,[10] the provisions set out below represent the core elements of this framework. They are presented as a concise reference for institutional leadership, governance bodies, and programme teams, ensuring that guidance is interpreted consistently and applied coherently across the sector. This framework consolidates existing Irish, European, and international standards within a single, values-based reference point, providing guidance on the responsible adoption of generative AI in ways consistent with the public mission of higher education in Ireland.

10 See https://hub.teachingandlearning.ie/genai/policy-framework

Principle 1: Academic Integrity, Transparency, and Accountability

Institutional AI Policy and Academic Freedom

Institutions should develop a single, coherent AI policy that defines permitted and prohibited uses across teaching and assessment. This policy should include discipline-sensitive exemplars, ensuring that academic freedom and pedagogical diversity are protected while maintaining consistency in standards and expectations. A clear institutional position fosters confidence among staff and students and supports alignment across the sector.

Transparency and Public Registers of Tools

AI systems required for student use should undergo formal institutional approval processes that include ethical review. Only those tools that are compliant with regulatory and ethical expectations should be adopted. Institutions are encouraged to maintain a publicly accessible register of approved tools, updated regularly with review criteria, risk assessments, safeguards, and retirement decisions. The purpose of this register is to demonstrate transparency and accountability. Given the evolving and imperfect nature of generative AI systems, institutions should communicate clearly about residual risks and mitigation measures, showing how decisions appropriately weigh educational benefit, ethical responsibility, and practical necessity.

Professional Development and Authentic Assessment

Effective integration of AI in higher education depends on well-supported staff and the continued renewal of assessment design. Institutions are encouraged to strengthen professional learning, peer exchange, and resource sharing that promote consistency and innovation in practice. Assessment approaches should be redesigned to prioritise authenticity, foregrounding student authorship and human judgment, as well as process-based learning. Programme design and workload planning should be aligned to ensure that these reforms are sustainable, equitable, and practicable.

AI Literacy Across the Curriculum

AI literacy should be embedded across programmes so that staff and students develop the capacity to critically evaluate and responsibly apply AI tools within their disciplines. Institutions may find it helpful to define learning outcomes that progress from foundational awareness to advanced critical engagement, ensuring that graduates can navigate AI-enhanced environments ethically and with informed judgment.

Oral Assessment Safeguard

To uphold fairness and academic integrity, institutions are advised to establish an institution-wide oral assessment safeguard that enables staff, regardless of any programme-level provisions or lack thereof, to demonstrate authorship directly, with the outcome of this process taking precedence over any existing written artefacts. Oral verification can help ensure authenticity without recourse to unreliable detection technologies. AI detectors and probabilistic tools should not be treated as determinative evidence of misconduct, and all integrity processes should rest on dialogue and evidence-based evaluation consistent with natural justice.

Principle 2: Equity & Inclusion

Institutional Commitment to Equity and Inclusion

Each institution should publish a clear statement of commitment to equity and inclusion within its AI policy, aligning with Irish equality law, the Public Sector Equality and Human Rights Duty, and the UN Sustainable Development Goal on inclusive education (SDG 4). This commitment signals that responsible AI use in higher education is inseparable from equality of opportunity and respect for diversity.

Embedding Equity in Institutional Practice

Equity considerations should be operationalised across procurement, staff and student development, and curriculum design. Institutions are encouraged to recognise that generative AI systems can reinforce or amplify existing and intersectional disadvantage, and to incorporate inclusion and accessibility into decision-making from the outset. Procurement and approval processes should therefore apply explicit equity criteria, seeking evidence of representative training data and awareness of system limitations. Where transparency or reliability cannot be demonstrated, institutions should adopt a precautionary approach and defer approval until identified risks can be appropriately mitigated.

Equitable Access to Tools and Infrastructure

Institutions should take steps to ensure that access to approved AI tools does not depend on students' private means. Shared or institutionally negotiated licensing arrangements support equitable access to approved AI tools. This principle extends to the digital infrastructure that enables AI use, including reliable broadband and hardware provision across disciplines, with capacity to support individual learners who may otherwise be excluded.

Linguistic Equity and Cultural Inclusion

Equity should also encompass linguistic and cultural diversity. Institutions should evaluate AI systems for their performance in the Irish language and, where necessary, provide appropriate supports or alternative arrangements.

Fairness in Assessment and Institutional Accountability

Assessment practices should reinforce, rather than undermine, equity. Institutions are encouraged to adopt the use of institutionally approved AI tools or require declarations when private systems are used, to safeguard fairness and comparability. Regular equity audits of AI adoption reporting disaggregated outcomes, providing clear complaint and redress pathways, and embedding student representation in AI governance structures, will help maintain accountability and ensure that commitments to inclusion translate into measurable action.

Principle 3: Critical Engagement, Human Oversight & AI Literacy

AI Literacy as a Core Graduate Attribute

Institutions are encouraged to embed AI literacy as a core graduate attribute across all programmes. This requires defining programme-specific learning outcomes that address technical foundations, disciplinary applications, ethical reasoning, and critical evaluation. A scaffolded progression from introductory awareness to advanced, discipline-specific engagement should be designed and assessed across the student journey. Such integration ensures that graduates can engage with gen AI critically and responsibly. This framework uses 'AI literacy' as its operative term, reflecting the current need to establish foundational competence across the sector. As generative AI becomes more deeply embedded in disciplinary practice and professional contexts, expectations may shift towards 'AI fluency', the capacity to work with these technologies as a routine and unremarkable part of intellectual and professional life. While AI literacy may be transdisciplinary, what fluency looks like will vary across disciplines, with fluent use of AI in the humanities differing markedly from fluent use in the sciences.

Professional Development and Interdisciplinary Engagement

Staff development is essential to credible AI literacy education. Institutions should provide structured opportunities for educators to develop confidence in teaching, assessing, and modelling responsible AI use within their disciplines. Interdisciplinary seminars or modules that bring together technical, ethical, and cultural perspectives can help staff and students situate gen AI within broader human contexts, reinforcing that critical engagement is as important as technical fluency.

Human Oversight and Accountability in Teaching and Assessment

Human oversight should remain a defining feature of all AI-enabled learning environments. Academic staff must retain final authority over assessment and curriculum decisions. Oversight expectations should also be embedded in procurement, including human-in-the-loop requirements and fitness-for-purpose declarations.

Institutional and Programme-Level Governance

Effective governance underpins ethical adoption. Institutions are encouraged to establish oversight mechanisms, such as committees or designated roles, with authority to review documentation, require bias testing, and where necessary, recommend the suspension of non-compliant tools. At local level, institutions are encouraged to identify clear points of responsibility to support staff, review practice, and raise concerns through existing governance channels.

Learning Pathways, Infrastructure, and Continuous Improvement

AI literacy development should be supported through coherent learning pathways, including consideration for mandatory induction for all students, optional advanced tracks for specialisation, and differentiated entry points that accommodate prior experience and accessibility needs. Institutions are encouraged to coordinate AI literacy initiatives across existing teaching and learning structures to maintain relevant curricula, exemplars, and responsive support for staff and students. Sustainable delivery depends on recognising the associated workload, ensuring access to suitable infrastructure, and embedding continuous improvement through transparent evaluation and student representation in governance.

Principle 4: Privacy & Data Governance

Transparency from Vendors and Institutions

Transparency should be treated as the cornerstone of responsible AI use. Institutions can promote transparency by seeking clear documentation from AI vendors before adoption. This documentation should set out the model's purpose, inputs and outputs, the provenance of training data, known limitations, and built-in safeguards. Such transparency enables informed institutional decision-making, supports compliance with legal and ethical obligations, and fosters public confidence in the use of AI for education. Institutions, in turn, should communicate with equal clarity. Plain-language transparency notices should be provided at the point of data collection and throughout processing, reflecting actual practices rather than generic templates. These notices should explain what data are collected, why, how they are used, and by whom, ensuring that staff and students can make informed choices about the use of specific gen AI systems.

Institutional Data Governance and Compliance

A coherent, institution-wide data governance framework is essential to guarantee compliance with GDPR, the EU AI Act, and national data-protection standards. Data collection should be strictly necessary for defined educational purposes, supported by a clear lawful basis and principles of proportionality. Data minimisation should be embedded within policy and practice, collecting and retaining only what is essential. Institutions are advised to document justifications for data use and to maintain visible accountability for compliance across governance structures.

Security and Risk Management

Institutions are responsible for maintaining robust security measures to protect personal and institutional data. Encryption, access controls, multi-factor authentication, and least-privilege access principles should be implemented as standard. Regular testing and independent security reviews help identify vulnerabilities and ensure timely remediation. Data Protection Impact Assessments (DPIAs) should be completed before any deployment or major system change and reviewed periodically to confirm that controls remain proportionate to risk and compliant with legal obligations.

Vendor and Contractual Accountability

Procurement processes should embed contractual protections that uphold institutional control and safeguard users' rights. Contracts with vendors should stipulate data ownership, prohibit training of external models without explicit consent, require enforceable data-deletion provisions, and include indemnities for breaches. Before approving any AI tool, institutions should seek sufficient transparency, documentation, and assurance of compliance to confirm that contractual terms align with ethical and legal standards.

Student Autonomy, Human Oversight, and Continuous Review

Students should retain authorship and intellectual-property rights over their work. No student data or content should be used for external model training without explicit, informed consent. Institutions are encouraged to ensure that high-stakes or consequential decisions involving generative AI always include meaningful human oversight and review. Regular evaluation of approved tools and transparency of oversight outcomes demonstrate institutional accountability and sustain public trust.

Principle 5: Sustainable Pedagogy

Sustainability as an Educational and Environmental Commitment

Institutions should integrate sustainability, both in terms of environmental responsibility and educational integrity, into all aspects of AI policy and practice. Sustainable adoption means not only reducing carbon and resource footprints but also ensuring that gen AI enhances rather than erodes the long-term quality of teaching and learning. Policies should make explicit how sustainability principles guide decision-making across procurement and infrastructure.

Environmental Impact and Vendor Accountability

Environmental considerations should be built into all gen AI procurement and adoption processes. Institutions are encouraged to conduct environmental impact assessments for proposed systems and to seek vendor disclosure on energy use and sustainability measures associated with their products. Preference should be given to efficient and low-energy models where these deliver comparable educational outcomes. Transparency on energy sourcing and carbon neutrality commitments should form part of procurement documentation.

Monitoring, Reporting, and Institutional Action

Institutions should monitor and review the environmental impact of gen AI use on an ongoing basis. Regular audits of AI-related energy consumption can provide evidence for targeted action to reduce ecological impact and inform wider institutional sustainability strategies. Findings should be actionable and, where possible, shared across the sector to support collective learning and continuous improvement.

Educational Sustainability and Capacity Building

Sustainable pedagogy depends on ongoing attention to both human and technological capacity. Institutions are encouraged to consider the long-term educational and operational implications of AI adoption alongside the continued development of teaching expertise. Responsible gen AI use should strengthen, not substitute, the pedagogical expertise and judgment that underpin higher education. Students can also be supported to understand the environmental impact of digital technologies through curricula that embed digital sustainability within broader AI literacy.

Resilience, Open Standards, and Continuous Review

To support institutional resilience, AI ecosystems should be designed to minimise dependence on any single platform or vendor. Procurement processes can incorporate preferences for open standards, data portability, and clear exit strategies to reduce the risk of vendor lock-in. Institutions are encouraged to review dependencies periodically to identify and address potential points of failure. Reviews of AI adoption can also consider environmental, educational, and financial sustainability in combination, ensuring that practice remains evidence-informed and transparent across the sector.

Monitoring and Sector Learning

Higher education institutions are encouraged to develop proportionate mechanisms for reflecting on how their engagement with generative AI aligns with this framework. These mechanisms should build on existing governance and quality assurance arrangements, ensuring that reflection and oversight are integrated rather than additional burdens.

The HEA will continue to monitor sector progress through established engagement processes and may invite institutions to share reflective accounts that address:

(1) Strategic decisions relating to generative AI in teaching, learning, and assessment;

(2) Developments in assessment design and academic integrity;

(3) Staff development and capacity-building initiatives that support AI literacy and pedagogical adaptation;

(4) Ethical, operational, or integrity concerns that have arisen and how they were addressed;

(5) Planning and investment decisions that enable responsible and sustainable AI adoption.

Institutions are encouraged to review generative AI developments regularly, given the rapid pace of technological change, and to share summaries of their reflections with staff, students, and the wider academic community. Doing so demonstrates a commitment to transparency and collective sector learning.

References

Analysis of Results from the Generative Artificial Intelligence Survey 2025. Quality and Qualifications Ireland, 2025. https://www.qqi.ie/what-we-do/engagement-insights-and-knowledge-sharing/our-data/analysis-of-generative-ai-survey.

Deng, Ruiqi, Maoli Jiang, Xinlu Yu, Yuyan Lu, and Shasha Liu. 'Does ChatGPT Enhance Student Learning? A Systematic Review and Meta-Analysis of Experimental Studies.' *Computers & Education* 227 (2025). https://doi.org/10.1016/j.compedu.2024.105224.

Elhussein, Gene, Elselot Hasselaar, Ostap Lutsyshyn, Tanya Milberg, and Saadia Zahidi. 'Shaping the Future of Learning: The Role of AI in Education 4.0.' World Economic Forum, 2024. https://www.weforum.org/publications/shaping-the-future-of-learning-the-role-of-ai-in-education-4-0/.

Ethical Guidelines on the Use of Artificial Intelligence (AI) and Data in Teaching and Learning for Educators. Publications Office of the European Union, 2022. https://data.europa.eu/doi/10.2766/153756.

Ethics Guidelines for Trustworthy AI. European Commission: Directorate-General for Communications Networks, Content and Technology and High-level Expert Group on Artificial Intelligence, 2019. https://doi.org/10.2759/346720.

Hemment, Drew, and Cody Kommers. 'Doing AI Differently.' The Alan Turing Institute, 2025. https://www.turing.ac.uk/news/publications/doing-ai-differently.

Hoernig, Steffen, André Ilharco, Paulo Trigo Pereira, and Regina Pereira. 'Generative AI and Higher Education: Challenges and Opportunities.' Institute of Public Policy, 2024. https://www.ipp-jcs.org/wp-content/uploads/2024/09/Report-AI-in-Higher-Education-IPP-1.pdf.

Lodge, Jason M., Sarah Howard, Margaret Bearman, and Phillip Dawson. 'Assessment Reform for the Age of Artificial Intelligence.' Tertiary Education Quality and Standards Agency, 2023. https://www.teqsa.gov.au/guides-resources/resources/corporate-publications/assessment-reform-age-artificial-intelligence.

Miao, Fengchun, and Wayne Holmes. 'Guidance for Generative AI in Education and Research.' UNESCO, 2023. https://doi.org/10.54675/EWZM9535.

Miao, Fengchun, Wayne Holmes, Ronghuai Huang, and Hui Zhang. 'AI and Education: Guidance for Policy-Makers.' UNESCO, 2021. https://doi.org/10.54675/PCSP7350.

Miao, Fengchun, and Cukurova Mutlu. 'AI Competency Framework for Teachers.' UNESCO, 2024. https://doi.org/10.54675/ZJTE2084.

Miao, Fengchun, Kelly Shiohira, and Natalie Leo. 'AI Competency Framework for Students.' UNESCO, 2024. https://doi.org/10.54675/JKJB9835.

O'Sullivan, James, Colin Lowry, Ross Woods, Brian Marrinan, and Clyde Hutchinson. *Generative AI in Higher Education Teaching and Learning: Sectoral Perspectives.* Higher Education Authority, 2025. https://zenodo.org/records/17153423.

Smeaton, Alan, Deirdre Ahern, Abeba Birhane, Susan Leavy, Bronagh Riordan, and Barry O'Sullivan. 'AI and Education.' Government of Ireland AI Advisory Council, 2025. https://www.gov.ie/en/department-of-enterprise-tourism-and-employment/publications/ai-advisory-council-advice-papers/.

Varsik, Samo, and Lydia Vosberg. 'The Potential Impact of Artificial Intelligence on Equity and Inclusion in Education.' OECD, 2024. https://doi.org/10.1787/15df715b-en.

Wang, Hui, Anh Dang, Zihao Wu, and Son Mac. 'Generative AI in Higher Education: Seeing ChatGPT through Universities' Policies, Resources, and Guidelines.' *Computers and Education: Artificial Intelligence* 7 (2024). https://doi.org/10.1016/j.caeai.2024.100326.

HEA | An tÚdarás um Ard-Oideachas
The Higher Education Authority

Higher Education Authority
3 Shelbourne Buildings
Crampton Avenue
Shelbourne Road
Ballsbridge
D04 C2Y6

Contact
Phone: +353 1 231 7100
Lo-Call Number: 1890 200 637
Fax: +353 1 231 7172
Email: info@hea.ie
www.hea.ie
𝕏 @hea_irl