

Generative AI in Higher Education Teaching & Learning

Evidence for the HEA's National Policy Framework

Contributors

James O'Sullivan
Colin Lowry
Ross Woods
Tim Conlon

HEA Generative AI Policy Framework

<https://hub.teachingandlearning.ie/genai/policy-framework>

HEA Generative AI Resource Portal

<https://hub.teachingandlearning.ie/genai/>

Version 1.0, December 2025


DOI: 10.82110/ph6k-mn37

Higher Education Authority, Dublin

How to cite:

O'Sullivan, James, Colin Lowry, Ross Woods & Tim Conlon. *Generative AI in Higher Education Teaching & Learning: Evidence for the HEA's National Policy Framework*. Dublin: Higher Education Authority, 2025. DOI: 10.82110/ph6k-mn37.

This document, and all original content contained within, is licensed under the [Creative Commons Attribution-ShareAlike 4.0 International Public License \(CC BY-SA 4.0\)](https://creativecommons.org/licenses/by-sa/4.0/).



Introduction

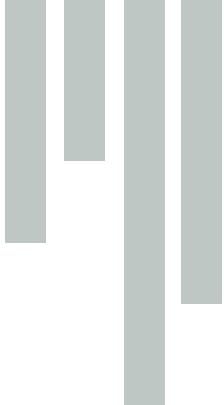
This evidence review underpins the Higher Education Authority's national policy framework on generative AI in teaching and learning. It draws on a substantial body of international policy, regulatory guidance, and research on the educational, ethical, and societal implications of generative systems. The objective is not to offer a snapshot of technological novelty, which would date quickly, but to identify durable principles, emerging empirical consensus, and the practical pressures now shaping institutional decision-making in Ireland.

International organisations and regulators have begun to stabilise a shared vocabulary for responsible AI. Work by UNESCO, the OECD and the European Commission has clarified what trustworthy systems should look like, who should be accountable for their use, and how human agency and equity must be protected when algorithmic tools enter sensitive public domains. These global reference points matter for Irish higher education, not least because universities are increasingly entangled with transnational platforms and procurement ecosystems, but because higher education has a public obligation to model ethical and evidence-informed engagement with new sociotechnical systems.

The practical integration of generative AI must be anchored in Ireland's regulatory environment, public service norms, linguistic and cultural commitments, and the distinctive mission of the sector. The Irish university is both a place of learning and a civic institution; it is also a public body expected to align with state principles governing data protection, transparency, procurement, and accountability. These constraints are not incidental and shape what responsible adoption can mean in classrooms, assessment design, institutional governance, and the student experience.

Generative AI is best understood here as an intensifier of long-running debates in higher education. Questions about authorship, originality, disciplinary method, academic labour, and assessment validity did not arrive with ChatGPT. What has changed is the pace at which students and staff can now produce polished outputs, and the extent to which generative capabilities are seeping into the routine software environment of academic work. This creates pressure on institutions to revisit assessment practices that over-reward product and under-evidence process, and to clarify what counts as appropriate assistance, partnership, or misconduct in an AI-saturated learning context.

Rather than rehearsing a simplistic opposition between innovation and risk, this review stays close to what the literature and sectoral consultation now suggest. The emerging direction is clear enough to justify firm policy steps in some areas, particularly around the limits of detection-led academic integrity



strategies, the need for assessment redesign, and the requirement for inclusive access and meaningful AI literacy. At the same time, the evidence base is still young in others, especially where claims about long-term cognitive effects, disciplinary variation, and the broader environmental footprint of AI systems remain contested or dependent on rapidly changing technical infrastructures.

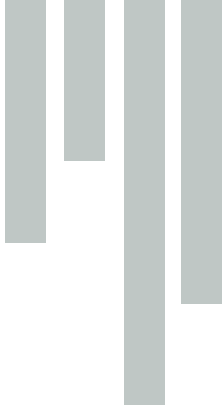
The sections that follow situate the HEA's policy work within this mixed landscape. They begin with the European ethical framework that continues to frame public-sector AI in Ireland, then address national guidance on AI in education and the public service, followed by the insights from the HEA's own focus-group consultations. The review then maps the international policy ecosystem and closes with the main peer-reviewed research trends in recent higher education literature.

Ethics Guidelines for Trustworthy AI

The European Commission High-Level Expert Group's *Ethics Guidelines for Trustworthy AI*¹ remain a foundational reference point for European AI governance. Published before the current wave of generative systems entered public use, the document nonetheless provides a durable account of what 'trustworthy AI' should mean in practice. It argues that trustworthiness rests on three interdependent conditions. AI should comply with law, align with ethical principles, and be technically robust enough to avoid predictable harm. While these and similar guidelines have been produced with developers in mind, they provide robust frameworks through which deployers and end-users, such as institutes of higher education, can assess the ethical and pedagogical consequences of specific AI systems and tools. Familiarity with such guidelines allows educators to ask to what extent a particular vendor attempted to make their system trustworthy, and where compromises have been made, what are the potential repercussions in terms of educational use.

The EC's guidelines identify four broad ethical principles: respect for human autonomy, prevention of harm, fairness, and explicability. These principles are translated into a practical set of requirements that can guide development and use across sectors. In this framing, AI should support human agency and remain subject to meaningful oversight; it should be safe, resilient, secure, and reliable; privacy and data governance should be treated as design fundamentals; system operations should be sufficiently

¹ *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>.



transparent and explainable for users to understand how outputs and decisions are produced; diversity and non-discrimination should shape both design and deployment; social and environmental impacts should be taken seriously rather than treated as externalities; accountability should be structured so that responsibility cannot be evaded when harms occur.

The guidelines also emphasise lifecycle thinking, that ethical safeguards must be embedded from problem definition through development, deployment, evaluation, and ongoing monitoring. They caution that tensions between requirements will create trade-offs that need to be made visible, justified, and documented. This insistence on explicit reasoning is particularly significant in educational contexts, where opaque systems can undermine trust even when intentions are benign.

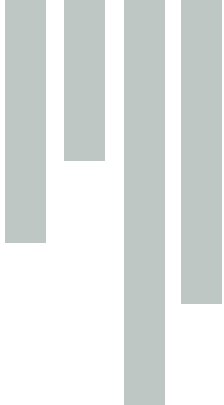
For higher education, the guidelines reinforce the primacy of academic judgement in teaching and assessment. They also set expectations for lawful and careful data handling in learning technologies, and underline the institutional responsibility to protect students from harm and to ensure that the pursuit of innovation does not erode fairness or the credibility of qualifications. The HEA's framework positions Irish higher education within a European commitment to responsible digital transformation, while also providing a principled basis for local policy choices.

AI and Education

The Government of Ireland's AI Advisory Council advice paper, *AI and Education*,² identifies generative AI as the most urgent and disruptive development currently facing Irish education. Although the report addresses the system as a whole rather than higher education alone, its analysis provides a valuable reminder that sector-specific responses must remain attentive to what is happening upstream and downstream. Students move across educational phases with habits, expectations, and inequalities that do not conveniently reset at the university gate.

The Council focuses deliberately on generative AI because it places powerful capabilities in the hands of teachers and learners without a commensurate level of public oversight. The tools are easy to access, difficult to regulate by prohibition alone, and already integrated into wider platforms that shape

² Alan Smeaton et al., '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/>.



everyday educational practice. The report acknowledges tangible opportunities for lesson planning, content generation, feedback, personalised explanation, and accessibility enhancements. It also notes that institutional responses remain uneven, with some settings attempting bans while others have encouraged experimentation or have left decisions to individuals. The resulting patchwork is a recipe for confusion, particularly around assessment expectations and academic integrity.

A notable contribution of the paper is its clear assessment of detection. The Council maintains that plagiarism detection tools cannot and will not provide a reliable solution to generative AI use. The policy implication is that institutions should invest in constructive integration and AI literacy rather than doubling down on technological enforcement that carries both technical and ethical weaknesses.

Equity sits at the centre of the Council's analysis, which highlights that subscription-based tools can widen disparities, while limited Irish-language support raises concerns about cultural and linguistic inclusion. The report therefore stresses that educational AI used by students should be private, secure, and free at the point of use, that data produced in educational contexts should not be repurposed for training and inclusion and accessibility must be designed into tool selection and deployment.

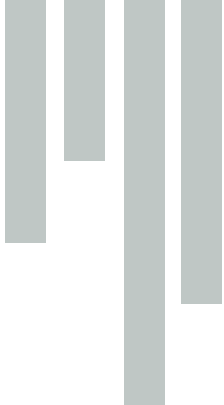
The Council recommends that national guidance should be treated as a living set of documents, updated in line with research and evolving capabilities. It also calls for sustained investment in AI literacy for educators. For higher education, these recommendations align with the HEA's consultation work and strengthen the case for coordinated national clarity that supports institutional autonomy without leaving practice to drift into incoherence.

Responsible Use of AI in the Public Service

The Department of Public Expenditure, NDP Delivery and Reform has issued the *Guidelines for the Responsible Use of Artificial Intelligence in the Public Service*.³ These guidelines adapt the European principles of trustworthy AI for the Irish state context and provide a practical benchmark for how public organisations should approach AI adoption.

The definition of responsible use offered here is grounded in three overarching commitments, that AI

³ 'Guidelines for the Responsible Use of AI in the Public Service,' Department of Public Expenditure, Infrastructure, Public Service Reform and Digitalisation, 2025, <https://gov.ie/en/department-of-public-expenditure-infrastructure-public-service-reform-and-digitalisation/publications/guidelines-for-the-responsible-use-of-ai-in-the-public-service/>.



should be people-centred, ethically governed, and demonstrably effective. The document reiterates the seven European requirements for trustworthiness and adds procedural expectations designed for the realities of public-sector decision-making.

The guidelines place particular weight on justification and proportionality. Public bodies are expected to articulate a clear public value rationale before adopting an AI system, including a consideration of whether non-AI approaches could meet the same goals. In terms of risk assessment, the likelihood and severity of harms must be weighed against expected benefits, with documented decisions and clear accountability.

Ethics, data protection, and oversight are expected to shape phases from problem definition and data collection to deployment, monitoring, and system retirement. Responsibility does not end once a tool is purchased or launched, and continuous evaluation, auditability, and redress mechanisms are expected when AI systems affect public outcomes.

Higher education institutions carry a dual obligation under this framework. They must meet these public service standards in administrative and educational uses of AI. They also have a wider civic role because they educate future professionals who will encounter AI across public and private spheres. When universities model careful, transparent adoption, they help build broader social expectations about how AI should be governed.

National focus groups

In April 2025, the Higher Education Authority (HEA) convened a comprehensive national consultation comprising ten themed focus groups and a leadership summit.⁴ Participants included institutional leaders, academic and professional staff, and representative student bodies from across Irish higher education. The purpose of these consultations was to move beyond abstract principle and capture the lived realities of how generative artificial intelligence is reshaping teaching, learning, and assessment. Their findings provide an essential evidence base for this framework, grounding national policy in the sector's own experience and ensuring alignment with both international standards and national

⁴ James O'Sullivan et al., *Generative AI in Higher Education Teaching and Learning: Sectoral Perspectives* (Higher Education Authority, 2025), <https://zenodo.org/records/17153423>.



priorities.

Participants consistently stressed that generative AI has not created new challenges for higher education but has intensified existing pressures: workload, student engagement, skills development, and the alignment of assessment with learning outcomes now appear more urgent. The rapid availability of AI tools has sharpened debates about what constitutes meaningful learning and fair assessment, compelling institutions to confront questions that had previously been postponed or addressed inconsistently.

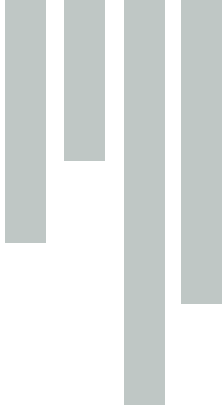
The focus groups revealed a diversity of perspectives. Some stakeholders regard generative AI as a threat to academic integrity and the credibility of qualifications, while others view it as an opportunity to innovate in assessment design and student support. Most agreed the reality lies between these poles: AI presents risks that must be carefully managed but also possibilities that should not be ignored. Absolutist positions—whether outright prohibition or uncritical adoption—were deemed unhelpful in practice.

Equity emerged as a recurring theme, and students reported uneven access to AI tools, often linked to the ability to pay for premium services or to navigate English-language interfaces. Staff highlighted that variations in institutional policy—from restrictive bans to permissive experimentation—create confusion for learners, particularly those moving across modules or programmes with inconsistent rules. The need for clarity, coherence, and consistency was strongly voiced.

The consultation underscored the importance of dialogue, with both staff and students alike calling for transparent institutional decision-making, greater opportunities to share practice, and clear communication of expectations in teaching and assessment. The sector demonstrated a strong willingness to engage constructively with AI, provided that policy development remains collaborative and responsive.

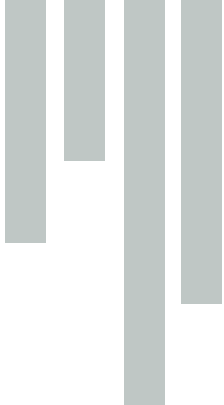
The HEA's analysis of these consultations distilled a set of system-wide findings to inform national and institutional planning:

1. Strategic coordination is urgently needed: Institutional responses remain uneven and reactive. A coordinated national approach is required to prevent fragmentation, maintain quality, and safeguard public trust.
2. Educational purpose requires re-articulation: Generative AI challenges established definitions of authorship, originality, and academic integrity. Core educational values must be reaffirmed and expressed for an AI-mediated context.

- 
3. Assessment reform is essential: Traditional text-based assignments are increasingly vulnerable to AI assistance. Assessment policy should prioritise authentic, process-focused, and AI-literate approaches over detection or enforcement.
 4. Equity and inclusion must be designed in: Without deliberate strategies, generative AI adoption may widen existing inequalities. Universal Design for Learning (UDL), inclusive digital-literacy initiatives, and equitable access to tools are critical.
 5. Capacity-building for staff and students is critical: Professional development, student partnership, and communities of practice are needed to build the pedagogical, ethical, and critical skills required for sustainable, values-led adoption.
 6. Governance and infrastructure must extend beyond technology: Institutions need clear ethical frameworks, transparent procurement, robust data-protection measures, and cross-functional leadership.
 7. Leadership must shape, not only manage, AI integration: Senior leaders are expected to articulate educational purpose and provide strategic direction, aligning national coherence with institutional autonomy.
 8. Dialogue and transparency are essential: Continuous, inclusive engagement with staff and students will foster trust, support shared norms, and ensure that policy remains adaptive as technologies evolve.

The consultation findings point to a series of strategic actions required to ensure that generative AI strengthens, rather than undermines, the mission of Irish higher education. These implications should guide both national policy development and institutional planning:

1. Values first: Keep human learning, critical thinking, and academic integrity at the centre of AI adoption. Policy and practice must reaffirm these principles as non-negotiable foundations of higher education.
2. Invest in people: Provide sustained funding for staff professional development, build AI literacy among students, and resource collaborative communities of practice so that educators and learners can engage critically and confidently with AI tools.
3. Update assessments: Support a sector-wide transition toward authentic, process-based assessment methods—such as oral examinations, staged submissions, and reflective portfolios—that prioritise originality, reasoning, and transparent use of AI where appropriate.
4. Embed inclusion: Apply Universal Design for Learning (UDL) principles and targeted digital-skills supports to ensure equitable access to AI tools and guard against the deepening of existing inequalities.
5. Strengthen governance: Establish shared national guidance while preserving institutional flexibility. Clear rules on acceptable AI use, robust data-privacy protections, and transparent procurement standards must underpin all adoption decisions.

- 
6. Foster continuous dialogue: Create formal mechanisms for ongoing consultation with staff, students, and sectoral partners to keep policy responsive to technological and pedagogical change.

These measures provide a roadmap for a values-led, evidence-informed, and sector-owned approach to generative AI, ensuring that innovation enhances educational integrity, equity, and the public trust in Irish higher education.

International policies and guidelines

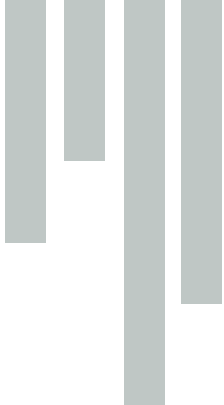
Across work issued by UNESCO, the OECD, the European Commission, various national regulators, the World Economic Forum, and a set of humanities-led interventions, gen AI is increasingly framed as a sociotechnical capacity that higher education must learn to shape rather than receive passively. Despite their different origins, these documents circle around a broadly coherent set of priorities: they place a human-centred ethical stance at the foundation of any educational use of AI, argue for a decisive move away from detection-driven approaches to integrity, call for more explicit competency development for both students and staff, and treat equity and strong institutional governance as conditions that must be secured rather than assumed.⁵

UNESCO places its guidance on gen AI within a longer normative tradition that treats education as a public endeavour grounded in dignity, agency, inclusion, and cultural-linguistic diversity.⁶ Its *Guidance for Generative AI in Education and Research* is explicit about the distance between rapid technological uptake and the slower pace of regulatory and institutional preparedness, drawing attention to data-protection risks, uneven capacity, and the absence of clear operational norms.⁷ It calls for transparency around how systems are deployed and sustained literacy work for both teachers and learners, alongside sector-specific rules that bring some coherence to institutional practice. Read together, these commitments tie classroom design to wider questions of governance, casting AI as an area in which public institutions retain responsibility for shaping the conditions under which learning takes place.

⁵ Drew Hemment and Cody Kommers, 'Doing AI Differently,' The Alan Turing Institute, 2025, <https://www.turing.ac.uk/news/publications/doing-ai-differently>.

⁶ Fengchun Miao et al., 'AI and Education: Guidance for Policy-Makers,' UNESCO, 2021, <https://doi.org/10.54675/PCSP7350>; Fengchun Miao and Wayne Holmes, 'Guidance for Generative AI in Education and Research,' UNESCO, 2023, <https://doi.org/10.54675/EWZM9535>.

⁷ Miao and Holmes, 'Guidance for Generative AI in Education and Research.'



The European Commission's *Ethical Guidelines on the Use of AI and Data in Teaching and Learning for Educators* operationalise trustworthy AI for staff through guiding questions, planning templates, and competence expectations aligned with EU data law and the (then-proposed) AI Act. The guidelines explicitly address data minimisation, purpose limitation, explainability, bias, and community engagement, making them a practical bridge between high-level ethics and day-to-day educational decisions.⁸

UNESCO's paired frameworks offer the most granular articulation of AI-related competences. The *AI Competency Framework for Teachers* defines competencies across human-centred mindset, ethics, AI foundations and applications, AI pedagogy, and AI for professional learning, organising these across progression levels intentionally framed as a master framework adaptable to local contexts. While based on slightly dated statistics, the framework notes a global capacity gap, with only a handful of countries having teacher AI frameworks or national CPD programmes.⁹ The *AI Competency Framework for Students* delineates similar competencies across human-centred mindset, ethics, AI techniques and applications, and AI system design. Its design positions learners as responsible co-creators able to scrutinise model outputs, reason about impacts, and engage in inclusive, sustainable design.¹⁰ Together, the frameworks embed human rights, inclusion, and agency, and they provide a scaffold for higher education curriculum mapping across levels, from foundation literacy to discipline-specific productive use, and, where appropriate, creation and stewardship.

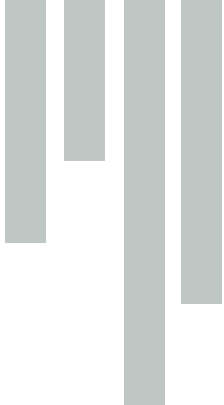
The TEQSA-commissioned discussion paper *Assessment Reform for the Age of Artificial Intelligence* synthesises expert consensus that non-invigilated tasks can rarely exclude generative AI use, detection is neither reliable nor ethically sufficient, and integrity must be reframed around appropriate engagement, evidence of learning processes, and capability development.¹¹ The paper sets out guiding principles and program-level propositions: diversify and sequence assessment across the degree, certify achievement at key points using targeted supervised assessments, teach and assess productive but ethical AI use explicitly, and build triangulated evidence of learning, such as process artefacts,

⁸ '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>.

⁹ Fengchun Miao and Cukurova Mutlu, 'AI Competency Framework for Teachers,' UNESCO, 2024, <https://doi.org/10.54675/ZJTE2084>.

¹⁰ Fengchun Miao et al., 'AI Competency Framework for Students,' UNESCO, 2024, <https://doi.org/10.54675/JKJB9835>.

¹¹ Jason M. Lodge et al., '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>.



prompt histories, oral defences, and interactive reviews. Illustrative redesigns show rebalancing of early-year exams toward formative coursework while (re)introducing supervised capstone-adjacent checks to assure standards.

These directions align with the WEF's *Shaping the Future of Learning*, which frames AI as augmenting rather than substituting teacher work. AI can relieve administrative load, support feedback-rich formative assessment, and personalise tasks, but the relational, motivational, and ethical labour of teaching remains irreducibly human.¹² The report also ties AI-enabled assessment to Education 4.0 goals, specifically global citizenship, innovation and creativity, technology fluency, and interpersonal skills.

Across international policy documents, a clear pattern has started to form. Most guidance highlights feedback and tutoring delivered at scale but always shaped by teacher-set parameters and supported by human oversight. It points to forms of content personalisation and accessibility support that widen participation while keeping educational judgement with staff, and it treats learner-analytics as useful only when indicators are interpreted rather than allowed to dictate decisions. It also encourages student production workflows that require prompts, sources, and revisions to be documented so that learning is evidenced through the movement of a task as well as through the final piece of work.¹³

European Commission guidelines stress educator agency, that staff must be able to explain why and how data are used, justify model-assisted decisions, and provide alternatives for students who opt out of automated processing.¹⁴

UNESCO's wider policy work makes the same connection between pedagogy and system design, urging institutions to teach students how to trace provenance, recognise model limitations, and understand the sources of bias that shape AI outputs. It also encourages the development of age-appropriate and context-specific use policies, coupled with sustained investment in teacher development, so that AI becomes part of the texture of learning design rather than an optional tool bolted on at the margins.¹⁵

The OECD provides a comprehensive equity lens across learner-facing, teacher-led, and institutional

¹² Gene Elhussein et al., '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/>.

¹³ Elhussein et al., '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.'

¹⁴ 'Ethical Guidelines on the Use of Artificial Intelligence (AI) and Data in Teaching and Learning for Educators.'

¹⁵ Miao et al., 'AI and Education: Guidance for Policy-Makers'; Miao and Holmes, 'Guidance for Generative AI in Education and Research.'



tools. Benefits (adaptivity, multimodal supports, early alerts) are paired with risks: infrastructural and affordability gaps, techno-ableism, bias amplification in admissions and risk models, privacy harms, and institutional capture by commercial platforms. The OECD recommends explicit equity impact assessments, culturally responsive design, transparent procurement and governance, and sustained professional learning so that staff can recognise and mitigate harms in situ.¹⁶

UNESCO's generative AI guidance centres inclusion and linguistic and cultural diversity, urging governments to evaluate long-term cognitive effects on critical thinking and creativity and to align AI-use expectations with realistic student practices.¹⁷ The European Commission likewise translates 'trustworthy AI' into classroom-level exigencies, namely, bias awareness, explainability, opt-outs, and community engagement.¹⁸

At the institutional level, the literature converges on whole-of-institution strategies that align governance, policy, procurement, digital infrastructure, staff development, academic regulation, and student partnership. UNESCO proposes integrated frameworks that connect AI adoption to inclusion, data governance, quality assurance, and monitoring.¹⁹ National and sectoral reports echo this framework. The Portuguese Institute of Public Policy's *Generative AI and Higher Education: Challenges and Opportunities* structures analysis by educators, students, and institutions, catalogues opportunities/threats, and issues targeted recommendations for risk minimisation while leveraging generative AI's potential.²⁰

Two cross-cutting governance tensions recur. First, assurance versus authenticity: institutions must assure standards without reverting to universal invigilation or surveillance, with programme-level architectures and triangulated evidence advanced as principled solutions.²¹ Second, balancing innovation with institutional risk, alongside procurement and data-sharing, requires due diligence.²²

¹⁶ Samo Varsik and Lydia Vosberg, 'The Potential Impact of Artificial Intelligence on Equity and Inclusion in Education,' OECD, 2024, <https://doi.org/10.1787/15df715b-en>.

¹⁷ Miao and Holmes, 'Guidance for Generative AI in Education and Research.'

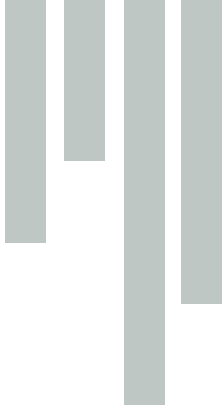
¹⁸ 'Ethical Guidelines on the Use of Artificial Intelligence (AI) and Data in Teaching and Learning for Educators.'

¹⁹ Miao et al., 'AI and Education: Guidance for Policy-Makers'; Miao and Holmes, 'Guidance for Generative AI in Education and Research.'

²⁰ Steffen Hoernig et al., '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>.

²¹ Lodge et al., 'Assessment Reform for the Age of Artificial Intelligence.'

²² 'Ethical Guidelines on the Use of Artificial Intelligence (AI) and Data in Teaching and Learning for Educators'; Varsik and Vosberg, 'The Potential Impact of Artificial Intelligence on Equity and Inclusion in Education.'



The Turing Institute's *Doing AI Differently* white paper argues that current systems exhibit a 'qualitative turn', generating and acting upon cultural outputs, such as language, images, and narratives, without adequate frameworks for contextual interpretation.²³ Two problems follow: homogenisation driven by narrow training data and benchmarking, and a deficit in interpretive judgement when ground truth is underdetermined. The paper calls for 'Interpretive AI' that embeds plurality, contestability, provenance, and situated reasoning into architectures and evaluation, and it explicitly positions universities, especially the humanities and arts, as co-designers of these capacities. For higher education, the implication is both curricular (training students to interrogate and shape AI systems) and infrastructural (creating public-interest labs and funding streams that reward interpretive depth, not only predictive accuracy).²⁴

In practical terms, the literature suggests that programme teams should treat generative AI as a design parameter within teaching and learning rather than a peripheral tool. Curricula should be mapped against staged AI competencies so that all students achieve baseline literacy, understanding model limits, provenance, and bias, while discipline-specific study develops productive, accountable use of AI for inquiry, analysis, and creation.²⁵ This curriculum work requires embedding opportunities for students to interrogate outputs, articulate their reasoning when they choose to use AI, and compare AI-enabled approaches with conventional methods. UNESCO's frameworks make this developmental arc explicit, positioning learners first as informed evaluators of AI, then as responsible users within their fields, and, where appropriate, as novice co-designers who can reason about data stewardship and system behaviour.²⁶

Assessment, in turn, must evidence learning processes as well as products. The consensus emerging from assessment scholarship and regulatory guidance is that detection-led integrity strategies are neither reliable nor pedagogically defensible, and assurance should instead come from redesigning tasks to elicit triangulated evidence of learning—planning notes, drafts, prompt histories, source triage, oral explanation, and reflective commentary—so that students 'show their work' when AI is present.²⁷ Programme-level provisions are particularly important: early stages can emphasise feedback-rich, low-stakes coursework that explicitly teaches appropriate AI use, while later stages reintroduce targeted

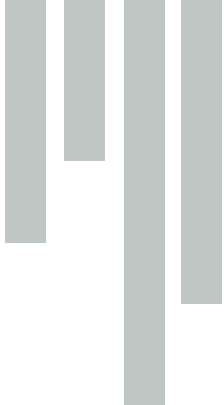
²³ Hemment and Kommers, 'Doing AI Differently.'

²⁴ Hemment and Kommers, 'Doing AI Differently.'

²⁵ Miao and Mutlu, 'AI Competency Framework for Teachers'; Miao et al., 'AI Competency Framework for Students.'

²⁶ Miao and Mutlu, 'AI Competency Framework for Teachers'; Miao et al., 'AI Competency Framework for Students.'

²⁷ Lodge et al., 'Assessment Reform for the Age of Artificial Intelligence.'



supervised assessments to certify capstone outcomes without defaulting to universal invigilation.²⁸ This approach aligns with Education 4.0 principles that favour continuous, formative feedback and authentic tasks over one-off, high-stakes examinations, provided that teachers retain interpretive judgment in calibrating and moderating AI-supported analytics.²⁹

Within teaching practice, AI can improve the cadence and quality of formative feedback, scaffold academic writing and problem solving, and widen accessibility through multimodal supports so long as its use is framed as augmentation and remains accountable to clearly articulated learning outcomes.³⁰ Effective integration therefore depends on assessment-ready rubrics that specify when and how AI may be used, what documentation students must provide, and how evaluators will weigh AI-mediated versus original contributions. These expectations should be taught, rather than simply policed, so that students learn to justify their choices, verify outputs against authoritative sources, and reflect on the intellectual trade-offs involved.³¹

Staff development is a precondition for any of the foregoing to work in classrooms. Educators require time and structured CPD to build fluency across three domains: foundational AI knowledge (capabilities, limits, and risks), pedagogy (designing activities and assessments that elicit evidence of learning in AI-rich settings), and professional practice (explaining decisions to students, documenting rationale, and iterating designs based on evaluation).³² When teachers co-design assessment criteria, trial AI-enabled feedback in controlled ways, and share exemplars within communities of practice, institutions see the strongest gains in learning quality without sacrificing standards.³³ A coherent strategy for teaching and assessment, anchored in transparent expectations, process evidence, and teacher judgment, allows higher education to harness whatever pedagogical benefits there may be in generative AI while maintaining academic integrity.

International guidance has moved beyond detection-versus-ban binaries toward a principled

²⁸ Lodge et al., 'Assessment Reform for the Age of Artificial Intelligence.'

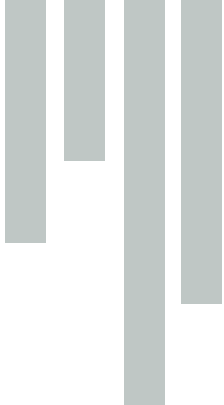
²⁹ Elhussein et al., 'Shaping the Future of Learning: The Role of AI in Education 4.0.'

³⁰ 'Ethical Guidelines on the Use of Artificial Intelligence (AI) and Data in Teaching and Learning for Educators'; Elhussein et al., 'Shaping the Future of Learning: The Role of AI in Education 4.0.'

³¹ 'Ethical Guidelines on the Use of Artificial Intelligence (AI) and Data in Teaching and Learning for Educators'; Miao and Holmes, 'Guidance for Generative AI in Education and Research.'

³² 'Ethical Guidelines on the Use of Artificial Intelligence (AI) and Data in Teaching and Learning for Educators'; Miao and Mutlu, 'AI Competency Framework for Teachers.'

³³ Lodge et al., 'Assessment Reform for the Age of Artificial Intelligence'; Miao and Holmes, 'Guidance for Generative AI in Education and Research.'




pragmatism, leveraging gen AI to expand access, create feedback-rich learning, and streamline administrative burdens within a rights-based regime that centres teacher agency, student development, and equity. UNESCO, the European Commission, the OECD, and TEQSA converge on program-level assessment redesign, explicit competency mapping, and coherent institutional governance. Extending this consensus, the Turing Institute argues for a deeper institutional mission with universities operating as co-designers of AI that can interpret context, sustain diversity, and keep human purposes in view. These disparate policy streams build a picture of an international landscape that is gradually settling on a shared direction of travel, in which generative AI is treated as a condition that reshapes how education is organised, governed, and justified, and in which the most credible responses rely on clarity of purpose, coherent institutional structures, and a willingness to redesign assessment and pedagogy so that learning processes are foregrounded rather than obscured by automated fluency. They also sketch a broader institutional role for higher education, one that asks universities to cultivate AI literacy, articulate defensible norms for ethical use, and participate in shaping systems that can sustain plurality, context, and human judgement, all while keeping equity and public accountability in view as the ground on which adoption should stand.

Research publications and position papers

Since ChatGPT was launched in late 2022, a recognisable pattern emerged across surveys, policy scans, and faculty reflections: students adopted generative AI quickly, staff experimented unevenly, and institutions wrote guidance faster than they could re-engineer their curricula. Quantitative snapshots and qualitative accounts converge on this storyline, with students describing generative AI as a useful tool for early-stage thinking and drafting, help with language and structure, and a way to reduce the ‘starting friction’ of complex tasks, while also registering unease about accuracy, dependence, and integrity.³⁴ Staff are, on average, more cautious, keen to try bounded uses, but held back by assessment

³⁴ Cecilia Ka Yuk Chan and Wenjie Hu, ‘Students’ Voices on Generative AI: Perceptions, Benefits, and Challenges in Higher Education,’ *International Journal of Educational Technology in Higher Education* 20 (2023), <https://doi.org/10.1186/s41239-023-00411-8>; David Baidoo-Anu et al., ‘Exploring Student Perspectives on Generative Artificial Intelligence in Higher Education Learning,’ *Discover Education* 3 (2024): 98, <https://doi.org/10.1007/s44217-024-00173-z>; Abayomi Arowosegbe et al., ‘Perception of Generative AI Use in UK Higher Education,’ *Frontiers in Education* 9 (2024), <https://doi.org/10.3389/feduc.2024.1463208>; Khoa Viet Nguyen, ‘The Use of Generative AI Tools in Higher Education: Ethical and Pedagogical Principles,’ *Journal of Academic Ethics* 23 (2025): 1435–55, <https://doi.org/10.1007/s10805-025->



risks, workload, and policy ambiguity.³⁵ Analyses of institutional guidance show a rapid shift from 'do not' to 'do, but disclose', and broad, principle-driven documents proliferated, often stronger on aspiration than on implementation.³⁶ The net effect is a policy–practice gap.

Within the last 12-18 months, the literature has moved beyond perceptions toward effects. Two meta-analyses synthesise dozens of experiments and quasi-experiments. Deng et al. estimate that ChatGPT assistance improves academic performance and reported higher-order thinking, reduces mental effort, and leaves self-efficacy mostly unchanged.³⁷ Wang and Fan analyse 51 studies and similarly report a large average effect on performance, with moderate gains in perceptions and higher-order outcomes; they also show that effect sizes vary with course type, learning model, and intervention duration.³⁸ These syntheses are careful about scope, noting that many primary studies are short, task-bounded, and language-adjacent (eg. writing, summarising, code explanation). But read together, they support a measured claim, that gen AI can reliably amplify near-term task performance in higher education, but generalisable learning depends on design.

Researchers argue that the field too often conflates assisted performance with unassisted

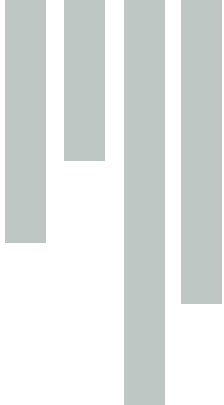
09607-1; Abdullahi Yusuf et al., 'Generative AI and the Future of Higher Education: A Threat to Academic Integrity or Reformation? Evidence from Multicultural Perspectives,' *International Journal of Educational Technology in Higher Education* 21 (2024), <https://doi.org/10.1186/s41239-024-00453-6>.

³⁵ Daniel Lee et al., 'The Impact of Generative AI on Higher Education Learning and Teaching: A Study of Educators' Perspectives,' *Computers and Education: Artificial Intelligence* 6 (2024), <https://doi.org/10.1016/j.caeai.2024.100221>; Armanto Sutedjo et al., 'Generative AI in Higher Education: A Cross-Institutional Study on Faculty Preparation and Resources,' *Studies in Technology Enhanced Learning* 4, no. 1 (2025), <https://doi.org/10.21428/8c225f6e.955a547e>; Dana-Kristin Mah et al., 'Perspectives of Academic Staff on Artificial Intelligence in Higher Education: Exploring Areas of Relevance,' *Frontiers in Education* 10 (2025), <https://doi.org/10.3389/feduc.2025.1484904>.

³⁶ Hui Wang et al., '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>; Nora McDonald et al., 'Generative Artificial Intelligence in Higher Education: Evidence from an Analysis of Institutional Policies and Guidelines,' *Computers in Human Behavior: Artificial Humans* 3 (2025), <https://doi.org/10.1016/j.chbah.2025.100121>; Attila Dabis and Csaba Csáki, 'AI and Ethics: Investigating the First Policy Responses of Higher Education Institutions to the Challenge of Generative AI,' *Humanities and Social Sciences Communications* 11 (2024), <https://doi.org/10.1057/s41599-024-03526-z>.

³⁷ Ruiqi Deng et al., '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>.

³⁸ Jin Wang and Wenxiang Fan, 'The Effect of ChatGPT on Students' Learning Performance, Learning Perception, and Higher-Order Thinking: Insights from a Meta-Analysis,' *Humanities and Social Sciences Communications* 12 (2025), <https://doi.org/10.1057/s41599-025-04787-y>.



competence.³⁹ If tasks assess end products that AI helps to polish, reported ‘learning gains’ may be artefacts of assistance rather than durable knowledge change. There are warnings that blanket offloading can attenuate metacognition and authorship unless instructors intentionally design for reflection, critique, and ‘fading’, that is, planned withdrawal of support.⁴⁰ A complementary strand—part empirical, part cultural commentary—frames this risk as cognitive debt, a build-up of unpractised sub-skills masked by fluent outputs. Evidence suggests reduced neural engagement during AI-assisted writing,⁴¹ giving rise to fears about stylistic convergence and voice flattening.⁴² The safest synthesis is not that generative AI harms learning, but that uncritical, persistent outsourcing of thinking sub-steps undermines the very difficulties that make learning stick.

Assessment practices shifted under pressure and not always coherently, and the technical literature now shows considerable agreement that AI-text detectors cannot carry the weight of high-stakes decisions because they misfire in both directions, can be bypassed with routine paraphrasing or translation, and display a persistent tendency to over-flag the work of non-native English writers.⁴³ Sector reporting mirrors the research, that institutions have experimented with detection and either reverted to invigilated exams or pivoted toward diversified assessment emphasising process evidence and oral verification. Across scoping reviews and practitioner frameworks, a clear pattern emerges in which redesigned tasks consistently outperform policing strategies, and programme teams are encouraged to use process portfolios, oral defences, authentic activities that tie work to context,


³⁹ Elisabeth Bauer et al., ‘Looking Beyond the Hype: Understanding the Effects of AI on Learning,’ *Educational Psychology Review* 37 (2025), <https://doi.org/10.1007/s10648-025-10020-8>; J. Weidlich et al., ‘ChatGPT in Education: An Effect in Search of a Cause,’ *Journal of Computer Assisted Learning* 41, no. 5 (2025), <https://doi.org/10.1111/jcal.70105>.

⁴⁰ Ali Darvishi et al., ‘Impact of AI Assistance on Student Agency,’ *Computers & Education* 210 (2024), <https://doi.org/10.1016/j.compedu.2023.104967>; Jasper Roe and Mike Perkins, ‘Generative AI and Agency in Education: A Critical Scoping Review and Thematic Analysis,’ preprint, arXiv, 2024, <https://doi.org/10.48550/arXiv.2411.00631>.

⁴¹ Nataliya Kosmyrna et al., ‘Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task,’ preprint, arXiv, 2025, <https://doi.org/10.48550/arXiv.2506.08872>.

⁴² Kyle Chayka, ‘A.I. Is Homogenizing Our Thoughts,’ *The New Yorker*, 2025, <https://www.newyorker.com/culture/infinite-scroll/ai-is-homogenizing-our-thoughts>.

⁴³ Debora Weber-Wulff et al., ‘Testing of Detection Tools for AI-Generated Text,’ *International Journal for Educational Integrity* 19 (2023), <https://doi.org/10.1007/s40979-023-00146-z>; Chaka Chaka, ‘Reviewing the Performance of AI Detection Tools in Differentiating between AI-Generated and Human-Written Texts: A Literature and Integrative Hybrid Review,’ *Journal of Applied Learning and Teaching* 7, no. 1 (2024): 115–26, <https://doi.org/10.37074/jalt.2024.7.1.14>; Jahna Otterbacher, ‘Why Technical Solutions for Detecting AI-Generated Content in Research and Education Are Insufficient,’ *Patterns* 4, no. 7 (2023), <https://doi.org/10.1016/j.patter.2023.100796>; Weixin Liang et al., ‘GPT Detectors Are Biased against Non-Native English Writers,’ *Patterns* 4, no. 7 (2023), <https://doi.org/10.1016/j.patter.2023.100779>.



transparent permission and disclosure regimes,⁴⁴ and structured opportunities to stress-test tasks that current generative models can already complete without difficulty.⁴⁵

The Irish context illustrates the signal-to-noise problem, as headlines about ‘AI cheating’ rose as cases were recorded, but the numbers remained small and categories blurry (eg. ‘plagiarism’ not separated from ‘AI-assisted misconduct’), confirming that counting incidents is a tricky governance strategy.⁴⁶ The better path, as argued by policy scholars and assessment specialists alike, is structural change that makes permitted use visible and meaningful, while raising the premium on explanation, transfer, and method.

Equity-centred studies and commentaries are strikingly consistent, and Addy et al. frame generative AI as an opportunity for transformative learning if institutions attend to access, literacies, and fair assessment cultures.⁴⁷ Work on accessibility documents concrete affordances for Universal Design for Learning while cautioning about hallucination risk and the need for human verification.⁴⁸ But the integrity layer can easily work against equity, because detectors tend to penalise multilingual writers and vague institutional rules leave students who depend on assistive technologies unsure whether the support they need is permitted or suspect, which places them in an avoidable and often stressful grey zone. Practitioner frameworks try to provide something firmer to stand on, using the RAFT structure—Rules, Access, Familiarity, Trust—and its later expansion to CRAFT to underline that effective practice depends as much on shared literacies and stable cultural norms as on the availability of tools, and that institutions need to build these conditions deliberately rather than hoping they will emerge on their

⁴⁴ Qi Xia et al., ‘A Scoping Review on How Generative Artificial Intelligence Transforms Assessment in Higher Education,’ *International Journal of Educational Technology in Higher Education* 21 (2024), <https://doi.org/10.1186/s41239-024-00468-z>; Jiahui Luo (Jess), ‘A Critical Review of GenAI Policies in Higher Education Assessment: A Call to Reconsider the ‘Originality’ of Students’ Work,’ *Assessment & Evaluation in Higher Education* 49, no. 5 (2024): 651–64; Mike Perkins et al., ‘The Artificial Intelligence Assessment Scale (AIAS): A Framework for Ethical Integration of Generative AI in Educational Assessment,’ *Journal of University Teaching and Learning Practice* 21, no. 06 (2024), <https://doi.org/10.53761/q3azde36>; Zachari Swiecki et al., ‘Assessment in the Age of Artificial Intelligence,’ *Computers and Education: Artificial Intelligence* 3 (2022), <https://doi.org/10.1016/j.caeai.2022.100075>.

⁴⁵ Binh Nguyen Thanh et al., ‘Race with the Machines: Assessing the Capability of Generative AI in Solving Authentic Assessments,’ *Australasian Journal of Educational Technology* 39, no. 5 (2023): 59–81, <https://doi.org/10.14742/ajet.8902>.

⁴⁶ Laura Lynott, ‘Level of Cheating at Irish Colleges Revealed – with AI Used in Many Instances of Plagiarism,’ *Irish Independent*, 2024, <https://www.independent.ie/irish-news/level-of-cheating-at-irish-colleges-revealed-with-ai-used-in-many-instances-of-plagiarism/a2012100671.html>.

⁴⁷ Tracie Addy et al., ‘Who Benefits and Who Is Excluded?: Transformative Learning, Equity, and Generative Artificial Intelligence,’ *Journal of Transformative Learning* 10, no. 2 (2023): 92–103.

⁴⁸ Anya S. Evmenova et al., ‘Harnessing the Power of Generative AI to Support ALL Learners,’ *TechTrends* 68 (2024): 820–31, <https://doi.org/10.1007/s11528-024-00966-x>.



own.⁴⁹

Research on learning design long predates the present hype cycle and still offers the most dependable guidance, since it treats technology as one element within a wider choreography of teaching rather than as a force that operates on its own. Within this literature, scholars use the idea of ‘orchestration’ to capture the way timing, role, and constraint shape educational value, and they show that when and how a tool is introduced into a task often matters far more than the specific features of the tool itself, because these design choices determine whether students remain actively engaged in reasoning or drift into passive outsourcing.⁵⁰ Translated into AI-era terms, this means making reasoning visible, regulating cognitive load, and allocating responsibility across humans and systems with an eye to agency.⁵¹ Darvishi et al. show that AI assistance modulates agency and effort,⁵² Steiss et al. find that model-generated writing feedback can approach novice-teacher quality on some dimensions *if* tethered to rubrics,⁵³ while Cordero, Torres-Zambrano, and Cordero-Castillo curate ‘best practices’ that are really design patterns.⁵⁴

The more ambitious personalisation promise remains debated. Generative AI can tailor tone and examples, but robust adaptation requires valid learner models and good data, which are not guaranteed in the wild.⁵⁵ That ambivalence appears inside creative-learning circles as well, and Resnick urges leveraging generative AI to broaden *creative* pathways while protecting the slow practices (tinkering, iteration, reflection) that foster originality,⁵⁶ while adjacent studies prototype AI-generated

⁴⁹ Danny Liu and Adam Bridgeman, ‘Rules, Access, Familiarity, and Trust – A Practical Approach to Addressing Generative AI in Education,’ *Teaching@Sydney*, 2024, <https://educational-innovation.sydney.edu.au/teaching@sydney/rules-access-familiarity-and-trust-a-practical-approach-to-addressing-generative-ai-in-education/>.

⁵⁰ Lucila Carvalho et al., ‘How Can We Design for Learning in an AI World?,’ *Computers and Education: Artificial Intelligence* 3 (2022), <https://doi.org/10.1016/j.caeai.2022.100053>; Hassan Khosravi et al., ‘Explainable Artificial Intelligence in Education,’ *Computers and Education: Artificial Intelligence* 3 (2022), <https://doi.org/10.1016/j.caeai.2022.100074>.

⁵¹ Dragan Gašević et al., ‘Empowering Learners for the Age of Artificial Intelligence,’ *Computers and Education: Artificial Intelligence* 4 (2023), <https://doi.org/10.1016/j.caeai.2023.100130>.

⁵² Darvishi et al., ‘Impact of AI Assistance on Student Agency.’

⁵³ Jacob Steiss et al., ‘Comparing the Quality of Human and ChatGPT Feedback of Students’ Writing,’ *Learning and Instruction* 91 (2024), <https://doi.org/10.1016/j.learninstruc.2024.101894>.

⁵⁴ Jorge Cordero et al., ‘Integration of Generative Artificial Intelligence in Higher Education: Best Practices,’ *Education Sciences* 15, no. 1 (2025), <https://doi.org/10.3390/educsci15010032>.

⁵⁵ Kristjan-Julius Laak et al., ‘Personalisation Is Not Guaranteed: The Challenges of Using Generative AI for Personalised Learning,’ in *Innovative Technologies and Learning*, ed. Yu-Ping Cheng et al. (Springer Nature Switzerland, 2024), https://doi.org/10.1007/978-3-031-65881-5_5; Ivica Pesovski et al., ‘Generative AI for Customizable Learning Experiences,’ *Sustainability* 16, no. 7 (2024), <https://doi.org/10.3390/su16073034>.

⁵⁶ Mitchel Resnick, ‘Generative AI and Creative Learning: Concerns, Opportunities, and Choices,’ *An MIT Exploration of Generative AI*, ahead of print, 2024, <https://doi.org/10.21428/e4baedd9.cf3e35e5>.



reflection prompts to scaffold self-directed learning.⁵⁷

Clear governance currents run in parallel: rapid policy production with uneven operationalisation,⁵⁸ rights-based proposals⁵⁹ that push universities beyond 'permitted uses' toward enforceable entitlements, and collective, critical stances urging democratic governance and labour recognition in response to the recognition that gen AI fundamentally re-shapes academic labour, epistemic authority, and student agency.⁶⁰ Practitioner synthesis tends to sit between these currents, listing near-term 'damage minimisation' moves, from slowing adoption where evidence is weak to strengthening assessment clarity and student autonomy,⁶¹ that is, what students and staff must be able to do in an AI-saturated university.

It is important to note that critiques of large language models in education are not purely abstract and Bender et al. warn that the race to scale models entrenches opacity, bias, and externalised environmental costs,⁶² with education-specific scoping reviews reaffirming these risks.⁶³ Energy and water use have been emphasised since the launch of ChatGPT, with some estimates suggesting a typical GPT-4o-class query uses on the order of 0.3 Wh, lower than early claims, though uncertainty remains and likely rises with more compute-hungry successors.⁶⁴ The broader point for campus strategy is that environmental externalities are real yet model- and workload-dependent, making

⁵⁷ Dishita Turakhia et al., 'Generating Reflection Prompts in Self-Directed Learning Activities with Generative AI,' *An MIT Exploration of Generative AI*, ahead of print, 2024, <https://doi.org/10.21428/e4baedd9.5970fe13>.

⁵⁸ Wang et al., 'Generative AI in Higher Education: Seeing ChatGPT through Universities' Policies, Resources, and Guidelines'; McDonald et al., 'Generative Artificial Intelligence in Higher Education: Evidence from an Analysis of Institutional Policies and Guidelines.'

⁵⁹ Kathryn Conrad, 'A Blueprint for an AI Bill of Rights for Education,' *Critical AI* 2, no. 1 (2024), <https://doi.org/10.1215/2834703X-11205245>.

⁶⁰ Aras Bozkurt et al., 'The Manifesto for Teaching and Learning in a Time of Generative AI: A Critical Collective Stance to Better Navigate the Future,' *Open Praxis* 16, no. 4 (2024), <https://doi.org/10.55982/openpraxis.16.4.777>.

⁶¹ Arran Hamilton et al., 'The Future of AI in Education: 13 Things We Can Do to Minimize the Damage,' preprint, OSF, 2023, <https://doi.org/10.35542/osf.io/372vr>.

⁶² Emily M. Bender et al., 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?,' *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Association for Computing Machinery, 2021, 610–23, <https://doi.org/10.1145/3442188.3445922>.

⁶³ Lixiang Yan et al., 'Practical and Ethical Challenges of Large Language Models in Education: A Systematic Scoping Review,' *British Journal of Educational Technology* 55, no. 1 (2024): 90–112, <https://doi.org/10.1111/bjet.13370>; Yifan Yao et al., 'A Survey on Large Language Model (LLM) Security and Privacy: The Good, The Bad, and The Ugly,' *High-Confidence Computing* 4, no. 2 (2024), <https://doi.org/10.1016/j.hcc.2024.100211>.

⁶⁴ Josh You, 'How Much Energy Does ChatGPT Use?,' Epoch AI, 2025, <https://epoch.ai/gradient-updates/how-much-energy-does-chatgpt-use>.



transparent reporting and demand management part of responsible adoption.

Security and authorship concerns also matter pedagogically, as LLMs are leaky reasoners with attack surfaces (prompt injection, data leakage), so ‘human-in-the-loop’ is a real assessment and safety requirement.⁶⁵ Critiques of anthropomorphic metaphors like ‘AI tutors’ and ‘AI collaborators’ warn that they obscure power and shift accountability away from institutions, with Sparrow and Flenady and Holmes et al. calling for political-economic analysis, not just classroom tactics.⁶⁶ Nowhere are these tensions starker than in the humanities, where writing is both method and product.

Stepping outside empirical studies, recent essays by Burnett, Hsu, and Underwood situate gen AI within older debates about the purpose of higher education.⁶⁷ Burnett’s classroom narrative—students conversing with a chatbot about ‘attention’ and then reflecting in prose—shows that AI can be a foil for human inquiry rather than its replacement, if teachers foreground the experience of thinking rather than the efficiency of answer-getting. Hsu’s field report from first-year writing argues that unreflective AI use corrodes process and voice, but also forces us to ask what the ‘English paper’ was for in the first place. Underwood suggests that LLMs have turned machine learning into ‘models of human culture’, creating a new, if uneasy, space for dialogue across computer science and humanistic method, provided we teach students to interrogate models, not just consume their outputs. Editorial reflection in *Nature Reviews Bioengineering* casts the point succinctly: writing is a way of thinking, so if gen AI shortcuts weaken the struggle to articulate, they weaken the learning.⁶⁸

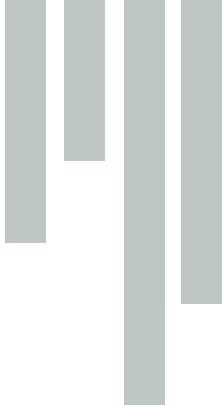
Across systematic reviews, design research, and critical commentary, a maturing consensus is visible, that generative AI can improve near-term performance but learning gains are conditional. Meta-analyses report positive average effects, moderated by task and design. Gains are most reliable where

⁶⁵ Yao et al., ‘A Survey on Large Language Model (LLM) Security and Privacy: The Good, The Bad, and The Ugly.’

⁶⁶ Robert Sparrow and Gene Flenady, ‘Bullshit Universities: The Future of Automated Education,’ *AI & Society*, ahead of print, 2025, <https://doi.org/10.1007/s00146-025-02340-8>; Wayne Holmes et al., ‘Critical Studies of Artificial Intelligence and Education: Putting a Stake in the Ground,’ preprint, Social Science Research Network, 2025, <https://doi.org/10.2139/ssrn.5391793>.

⁶⁷ D. Graham Burnett, ‘Will the Humanities Survive Artificial Intelligence?’, *The New Yorker*, 2025, <https://www.newyorker.com/culture/the-weekend-essay/will-the-humanities-survive-artificial-intelligence>; Hua Hsu, ‘What Happens After A.I. Destroys College Writing?’, *The New Yorker*, 2025, <https://www.newyorker.com/magazine/2025/07/07/the-end-of-the-english-paper>; Ted Underwood, ‘The Impact of Language Models on the Humanities and Vice Versa,’ *Nature Computational Science* 5 (2025), <https://doi.org/10.1038/s43588-025-00819-4>.

⁶⁸ ‘Writing Is Thinking,’ *Nature Reviews Bioengineering* 3 (2025), <https://doi.org/10.1038/s44222-025-00323-4>.



AI support is embedded in well-scaffolded activities that require critique, transformation, and explanation.⁶⁹ Detectors are unreliable and inequitable, so assessment must change structurally. The sector is moving toward tasks that surface process and reasoning, with transparent permission and disclosure and viva-style verification where appropriate.⁷⁰

Equity is hard won, and while generative AI can reduce barriers (language, accessibility), policy choices and resource gaps can re-stratify advantage. Frameworks like RAFT/CRAFT point to what institutions need to embed in culture in terms of rules, access, familiarity, and trust.⁷¹ Governance should be rights- and capability-based.

The literature traces a field that is moving steadily from early fascination and scattered experimentation toward a more grounded understanding of what generative AI can and cannot offer higher education, and it shows that the most durable insights come from studies that attend to design, process, and equity rather than to novelty or scale. The evidence now points to a simple but demanding conclusion, that near-term gains in performance are achievable but only become meaningful when accompanied by tasks that reveal reasoning and encourage students to test, question, and justify their use of AI, and that institutions need to create environments in which such practices can take root. Across the research, the same priorities recur: students must be able to show how they learn rather than simply present polished artefacts, staff need time and support to redesign assessment and pedagogy, and governance has to be built around rights, capabilities, transparency, and fair access rather than around restrictive policing. What emerges is not a call for either enthusiasm or alarm, but a recognition that responsible adoption requires attention to culture as much as to tools, and that universities will need to cultivate shared norms, clear expectations, and coherent structures if generative AI is to enhance rather than erode the conditions under which learning and academic judgement remain credible.

⁶⁹ Deng et al., 'Does ChatGPT Enhance Student Learning? A Systematic Review and Meta-Analysis of Experimental Studies'; Wang and Fan, 'The Effect of ChatGPT on Students' Learning Performance, Learning Perception, and Higher-Order Thinking: Insights from a Meta-Analysis.'

⁷⁰ Weber-Wulff et al., 'Testing of Detection Tools for AI-Generated Text'; Liang et al., 'GPT Detectors Are Biased against Non-Native English Writers'; Xia et al., 'A Scoping Review on How Generative Artificial Intelligence Transforms Assessment in Higher Education'; Luo (Jess), 'A Critical Review of GenAI Policies in Higher Education Assessment.'

⁷¹ Liu and Bridgeman, 'Rules, Access, Familiarity, and Trust – A Practical Approach to Addressing Generative AI in Education'; Addy et al., 'Who Benefits and Who Is Excluded?: Transformative Learning, Equity, and Generative Artificial Intelligence.'


References

- Addy, Tracie, Tingting Kang, Tim Laquintano, and Vivienne Dietrich. 2023. 'Who Benefits and Who Is Excluded?: Transformative Learning, Equity, and Generative Artificial Intelligence.' *Journal of Transformative Learning* 10 (2): 92–103.
- Arowosegbe, Abayomi, Jaber S. Alqahtani, and Tope Oyelade. 2024. 'Perception of Generative AI Use in UK Higher Education.' *Frontiers in Education* 9. <https://doi.org/10.3389/feduc.2024.1463208>.
- Baidoo-Anu, David, Daniel Asamoah, Isaac Amoako, and Inuusah Mahama. 2024. 'Exploring Student Perspectives on Generative Artificial Intelligence in Higher Education Learning.' *Discover Education* 3: 98. <https://doi.org/10.1007/s44217-024-00173-z>.
- Bauer, Elisabeth, Samuel Greiff, Arthur C. Graesser, Katharina Scheiter, and Michael Sailer. 2025. 'Looking Beyond the Hype: Understanding the Effects of AI on Learning.' *Educational Psychology Review* 37. <https://doi.org/10.1007/s10648-025-10020-8>.
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?' *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–23. <https://doi.org/10.1145/3442188.3445922>.
- Bozkurt, Aras, Junhong Xiao, Robert Farrow, et al. 2024. 'The Manifesto for Teaching and Learning in a Time of Generative AI: A Critical Collective Stance to Better Navigate the Future.' *Open Praxis* 16 (4). <https://doi.org/10.55982/openpraxis.16.4.777>.
- Burnett, D. Graham. 2025. 'Will the Humanities Survive Artificial Intelligence?' *The New Yorker*. <https://www.newyorker.com/culture/the-weekend-essay/will-the-humanities-survive-artificial-intelligence>.
- Carvalho, Lucila, Roberto Martinez-Maldonado, Yi-Shan Tsai, Lina Markauskaite, and Maarten De Laat. 2022. 'How Can We Design for Learning in an AI World?' *Computers and Education: Artificial Intelligence* 3. <https://doi.org/10.1016/j.caeai.2022.100053>.
- Chaka, Chaka. 2024. 'Reviewing the Performance of AI Detection Tools in Differentiating between AI-Generated and Human-Written Texts: A Literature and Integrative Hybrid Review.' *Journal of Applied Learning and Teaching* 7 (1): 115–26. <https://doi.org/10.37074/jalt.2024.7.1.14>.
- Chan, Cecilia Ka Yuk, and Wenjie Hu. 2023. 'Students' Voices on Generative AI: Perceptions, Benefits, and Challenges in Higher Education.' *International Journal of Educational Technology in Higher Education* 20. <https://doi.org/10.1186/s41239-023-00411-8>.
- Chayka, Kyle. 2025. 'A.I. Is Homogenizing Our Thoughts.' *The New Yorker*. <https://www.newyorker.com/culture/infinite-scroll/ai-is-homogenizing-our-thoughts>.
- Conrad, Kathryn. 2024. 'A Blueprint for an AI Bill of Rights for Education.' *Critical AI* 2 (1). <https://doi.org/10.1215/2834703X-11205245>.
- Cordero, Jorge, Jonathan Torres-Zambrano, and Alison Cordero-Castillo. 2025. 'Integration of

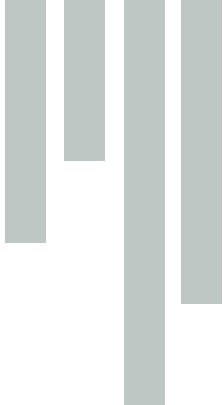
- Generative Artificial Intelligence in Higher Education: Best Practices.' *Education Sciences* 15 (1). <https://doi.org/10.3390/educsci15010032>.
- Dabis, Attila, and Csaba Csáki. 2024. 'AI and Ethics: Investigating the First Policy Responses of Higher Education Institutions to the Challenge of Generative AI.' *Humanities and Social Sciences Communications* 11. <https://doi.org/10.1057/s41599-024-03526-z>.
- Darvishi, Ali, Hassan Khosravi, Shazia Sadiq, Dragan Gašević, and George Siemens. 2024. 'Impact of AI Assistance on Student Agency.' *Computers & Education* 210. <https://doi.org/10.1016/j.compedu.2023.104967>.
- Deng, Ruiqi, Maoli Jiang, Xinlu Yu, Yuyan Lu, and Shasha Liu. 2025. 'Does ChatGPT Enhance Student Learning? A Systematic Review and Meta-Analysis of Experimental Studies.' *Computers & Education* 227. <https://doi.org/10.1016/j.compedu.2024.105224>.
- Elhussein, Gene, Elseot Hasselaar, Ostap Lutsyshyn, Tanya Milberg, and Saadia Zahidi. 2024. 'Shaping the Future of Learning: The Role of AI in Education 4.0.' World Economic Forum. <https://www.weforum.org/publications/shaping-the-future-of-learning-the-role-of-ai-in-education-4-0/>.
- 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>.
- 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>.
- Evmenova, Anya S., Jered Borup, and Joan Kang Shin. 2024. 'Harnessing the Power of Generative AI to Support ALL Learners.' *TechTrends* 68: 820–31. <https://doi.org/10.1007/s11528-024-00966-x>.
- Gašević, Dragan, George Siemens, and Shazia Sadiq. 2023. 'Empowering Learners for the Age of Artificial Intelligence.' *Computers and Education: Artificial Intelligence* 4. <https://doi.org/10.1016/j.caeai.2023.100130>.
- Guidelines for the Responsible Use of AI in the Public Service*. 2025. Department of Public Expenditure, Infrastructure, Public Service Reform and Digitalisation. <https://gov.ie/en/departments-of-public-expenditure-infrastructure-public-service-reform-and-digitalisation/publications/guidelines-for-the-responsible-use-of-ai-in-the-public-service/>.
- Hamilton, Arran, Dylan Wiliam, and John Hattie. 2023. 'The Future of AI in Education: 13 Things We Can Do to Minimize the Damage.' Preprint, OSF. <https://doi.org/10.35542/osf.io/372vr>.
- Hemment, Drew, and Cody Kommers. 2025. 'Doing AI Differently.' The Alan Turing Institute. <https://www.turing.ac.uk/news/publications/doing-ai-differently>.
- Hoernig, Steffen, André Ilharco, Paulo Trigo Pereira, and Regina Pereira. 2024. 'Generative AI and Higher Education: Challenges and Opportunities.' Institute of Public Policy. <https://www.ipp-jcs.org/wp-content/uploads/2024/09/Report-AI-in-Higher-Education-IPP-1.pdf>.
- Holmes, Wayne, Ana Mouta, Velislava Hillman, et al. 2025. 'Critical Studies of Artificial Intelligence and Education: Putting a Stake in the Ground.' Preprint, Social Science Research Network.

- <https://doi.org/10.2139/ssrn.5391793>.
- Hsu, Hua. 2025. 'What Happens After A.I. Destroys College Writing?' *The New Yorker*.
<https://www.newyorker.com/magazine/2025/07/07/the-end-of-the-english-paper>.
- Khosravi, Hassan, Simon Buckingham Shum, Guanliang Chen, et al. 2022. 'Explainable Artificial Intelligence in Education.' *Computers and Education: Artificial Intelligence* 3.
<https://doi.org/10.1016/j.caeai.2022.100074>.
- Kosmyna, Nataliya, Eugene Hauptmann, Ye Tong Yuan, et al. 2025. 'Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task.' Preprint, arXiv. <https://doi.org/10.48550/arXiv.2506.08872>.
- Laak, Kristjan-Julius, Rania Abdelghani, and Jaan Aru. 2024. 'Personalisation Is Not Guaranteed: The Challenges of Using Generative AI for Personalised Learning.' In *Innovative Technologies and Learning*, edited by Yu-Ping Cheng, Margus Pedaste, Emanuele Bardone, and Yueh-Min Huang. Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-65881-5_5.
- Lee, Daniel, Matthew Arnold, Amit Srivastava, et al. 2024. 'The Impact of Generative AI on Higher Education Learning and Teaching: A Study of Educators' Perspectives.' *Computers and Education: Artificial Intelligence* 6. <https://doi.org/10.1016/j.caeai.2024.100221>.
- Liang, Weixin, Mert Yuksekgonul, Yining Mao, Eric Wu, and James Zou. 2023. 'GPT Detectors Are Biased against Non-Native English Writers.' *Patterns* 4 (7). <https://doi.org/10.1016/j.patter.2023.100779>.
- Liu, Danny, and Adam Bridgeman. 2024. 'Rules, Access, Familiarity, and Trust – A Practical Approach to Addressing Generative AI in Education.' *Teaching@Sydney*. <https://educational-innovation.sydney.edu.au/teaching@sydney/rules-access-familiarity-and-trust-a-practical-approach-to-addressing-generative-ai-in-education/>.
- Lodge, Jason M., Sarah Howard, Margaret Bearman, and Phillip Dawson. 2023. 'Assessment Reform for the Age of Artificial Intelligence.' Tertiary Education Quality and Standards Agency. <https://www.teqsa.gov.au/guides-resources/resources/corporate-publications/assessment-reform-age-artificial-intelligence>.
- Luo (Jess), Jiahui. 2024. 'A Critical Review of GenAI Policies in Higher Education Assessment: A Call to Reconsider the 'Originality' of Students' Work.' *Assessment & Evaluation in Higher Education* 49 (5): 651–64.
- Lynott, Laura. 2024. 'Level of Cheating at Irish Colleges Revealed – with AI Used in Many Instances of Plagiarism.' *Irish Independent*. <https://www.independent.ie/irish-news/level-of-cheating-at-irish-colleges-revealed-with-ai-used-in-many-instances-of-plagiarism/a2012100671.html>.
- Mah, Dana-Kristin, Nils Knoth, and Marc Egloffstein. 2025. 'Perspectives of Academic Staff on Artificial Intelligence in Higher Education: Exploring Areas of Relevance.' *Frontiers in Education* 10. <https://doi.org/10.3389/feduc.2025.1484904>.
- McDonald, Nora, Aditya Johri, Areej Ali, and Aayushi Hingle Collier. 2025. 'Generative Artificial Intelligence in Higher Education: Evidence from an Analysis of Institutional Policies and Guidelines.' *Computers in Human Behavior: Artificial Humans* 3.

- 
- <https://doi.org/10.1016/j.chbah.2025.100121>.
- Miao, Fengchun, and Wayne Holmes. 2023. 'Guidance for Generative AI in Education and Research.' UNESCO. <https://doi.org/10.54675/EWZM9535>.
- Miao, Fengchun, Wayne Holmes, Ronghuai Huang, and Hui Zhang. 2021. 'AI and Education: Guidance for Policy-Makers.' UNESCO. <https://doi.org/10.54675/PCSP7350>.
- Miao, Fengchun, and Cukurova Mutlu. 2024. 'AI Competency Framework for Teachers.' UNESCO. <https://doi.org/10.54675/ZJTE2084>.
- Miao, Fengchun, Kelly Shiohira, and Natalie Leo. 2024. 'AI Competency Framework for Students.' UNESCO. <https://doi.org/10.54675/JKJB9835>.
- Nguyen, Khoa Viet. 2025. 'The Use of Generative AI Tools in Higher Education: Ethical and Pedagogical Principles.' *Journal of Academic Ethics* 23: 1435–55. <https://doi.org/10.1007/s10805-025-09607-1>.
- 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>.
- Otterbacher, Jahna. 2023. 'Why Technical Solutions for Detecting AI-Generated Content in Research and Education Are Insufficient.' *Patterns* 4 (7). <https://doi.org/10.1016/j.patter.2023.100796>.
- Perkins, Mike, Leon Furze, Jasper Roe, and Jason MacVaugh. 2024. 'The Artificial Intelligence Assessment Scale (AIAS): A Framework for Ethical Integration of Generative AI in Educational Assessment.' *Journal of University Teaching and Learning Practice* 21 (06). <https://doi.org/10.53761/q3azde36>.
- Pesovski, Ivica, Ricardo Santos, Roberto Henriques, and Vladimir Trajkovic. 2024. 'Generative AI for Customizable Learning Experiences.' *Sustainability* 16 (7). <https://doi.org/10.3390/su16073034>.
- Resnick, Mitchel. 2024. 'Generative AI and Creative Learning: Concerns, Opportunities, and Choices.' *MIT Exploration of Generative AI*, ahead of print. <https://doi.org/10.21428/e4baedd9.cf3e35e5>.
- Roe, Jasper, and Mike Perkins. 2024. 'Generative AI and Agency in Education: A Critical Scoping Review and Thematic Analysis.' Preprint, arXiv. <https://doi.org/10.48550/arXiv.2411.00631>.
- Smeaton, Alan, Deirdre Ahern, Abeba Birhane, Susan Leavy, Bronagh Riordan, and Barry O'Sullivan. 2025. 'AI and Education.' Government of Ireland AI Advisory Council. <https://www.gov.ie/en/department-of-enterprise-tourism-and-employment/publications/ai-advisory-council-advice-papers/>.
- Sparrow, Robert, and Gene Flenady. 2025. 'Bullshit Universities: The Future of Automated Education.' *AI & Society*, ahead of print. <https://doi.org/10.1007/s00146-025-02340-8>.
- Steiss, Jacob, Tamara Tate, Steve Graham, et al. 2024. 'Comparing the Quality of Human and ChatGPT Feedback of Students' Writing.' *Learning and Instruction* 91. <https://doi.org/10.1016/j.learninstruc.2024.101894>.
- Sutedjo, Armanto, Stephanie Peitseng Liu, and Mahjabin Chowdhury. 2025. 'Generative AI in Higher Education: A Cross-Institutional Study on Faculty Preparation and Resources.' *Studies in Technology Enhanced Learning* 4 (1). <https://doi.org/10.21428/8c225f6e.955a547e>.
- Swiecki, Zachari, Hassan Khosravi, Guanliang Chen, et al. 2022. 'Assessment in the Age of Artificial



- Intelligence.' *Computers and Education: Artificial Intelligence* 3. <https://doi.org/10.1016/j.caeai.2022.100075>.
- Thanh, Binh Nguyen, Diem Thi Hong Vo, Minh Nguyen Nhat, Thi Thu Tra Pham, Hieu Thai Trung, and Son Ha Xuan. 2023. 'Race with the Machines: Assessing the Capability of Generative AI in Solving Authentic Assessments.' *Australasian Journal of Educational Technology* 39 (5): 59–81. <https://doi.org/10.14742/ajet.8902>.
- Turakhia, Dishita, Zane Mroue, Peiling Jiang, and Stefanie Mueller. 2024. 'Generating Reflection Prompts in Self-Directed Learning Activities with Generative AI.' *An MIT Exploration of Generative AI*, ahead of print. <https://doi.org/10.21428/e4baedd9.5970fe13>.
- Underwood, Ted. 2025. 'The Impact of Language Models on the Humanities and Vice Versa.' *Nature Computational Science* 5. <https://doi.org/10.1038/s43588-025-00819-4>.
- Varsik, Samo, and Lydia Vosberg. 2024. 'The Potential Impact of Artificial Intelligence on Equity and Inclusion in Education.' OECD. <https://doi.org/10.1787/15df715b-en>.
- Wang, Hui, Anh Dang, Zihao Wu, and Son Mac. 2024. 'Generative AI in Higher Education: Seeing ChatGPT through Universities' Policies, Resources, and Guidelines.' *Computers and Education: Artificial Intelligence* 7. <https://doi.org/10.1016/j.caeai.2024.100326>.
- Wang, Jin, and Wenxiang Fan. 2025. 'The Effect of ChatGPT on Students' Learning Performance, Learning Perception, and Higher-Order Thinking: Insights from a Meta-Analysis.' *Humanities and Social Sciences Communications* 12. <https://doi.org/10.1057/s41599-025-04787-y>.
- Weber-Wulff, Debora, Alla Anohina-Naumeca, Sonja Bjelobaba, et al. 2023. 'Testing of Detection Tools for AI-Generated Text.' *International Journal for Educational Integrity* 19. <https://doi.org/10.1007/s40979-023-00146-z>.
- Weidlich, J., D. Gašević, H. Drachsler, and P. Kirschner. 2025. 'ChatGPT in Education: An Effect in Search of a Cause.' *Journal of Computer Assisted Learning* 41 (5). <https://doi.org/10.1111/jcal.70105>.
- 'Writing Is Thinking.' 2025. *Nature Reviews Bioengineering* 3. <https://doi.org/10.1038/s44222-025-00323-4>.
- Xia, Qi, Xiaojing Weng, Fan Ouyang, Tzung Jin Lin, and Thomas K.F. Chiu. 2024. 'A Scoping Review on How Generative Artificial Intelligence Transforms Assessment in Higher Education.' *International Journal of Educational Technology in Higher Education* 21. <https://doi.org/10.1186/s41239-024-00468-z>.
- Yan, Lixiang, Lele Sha, Linxuan Zhao, et al. 2024. 'Practical and Ethical Challenges of Large Language Models in Education: A Systematic Scoping Review.' *British Journal of Educational Technology* 55 (1): 90–112. <https://doi.org/10.1111/bjet.13370>.
- Yao, Yifan, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Zhibo Sun, and Yue Zhang. 2024. 'A Survey on Large Language Model (LLM) Security and Privacy: The Good, The Bad, and The Ugly.' *High-Confidence Computing* 4 (2). <https://doi.org/10.1016/j.hcc.2024.100211>.
- You, Josh. 2025. 'How Much Energy Does ChatGPT Use?' Epoch AI. <https://epoch.ai/gradient-updates/how-much-energy-does-chatgpt-use>.
- Yusuf, Abdullahi, Nasrin Pervin, and Marcos Román-González. 2024. 'Generative AI and the Future of



Higher Education: A Threat to Academic Integrity or Reformation? Evidence from Multicultural Perspectives.' *International Journal of Educational Technology in Higher Education* 21.
<https://doi.org/10.1186/s41239-024-00453-6>.