

Generative AI in Higher Education Teaching & Learning

Alignment with the EU AI Act

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

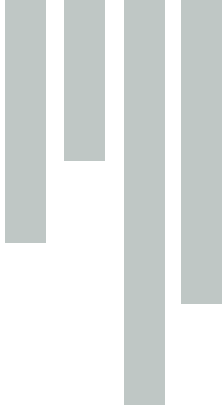
Higher Education Authority, Dublin

DOI: 10.82110/tgec-fq82

How to cite:

O'Sullivan, James, Colin Lowry, Ross Woods & Tim Conlon. *Generative AI in Higher Education Teaching & Learning: Alignment with the EU AI Act*. Dublin: Higher Education Authority, 2025. DOI: 10.82110/tgec-fq82.

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/).



This document is intended as policy guidance to support Irish higher education institutions in understanding and aligning teaching and learning practices with the EU AI Act. It does not constitute legal advice. Institutions should seek independent legal counsel when making compliance determinations or procurement decisions.

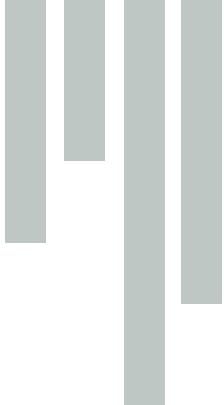
The European Union Artificial Intelligence Act (Regulation (EU) 2024/1689) establishes the first binding, cross-sectoral regulatory framework for artificial intelligence within the EU. As a regulation, it introduces enforceable obligations for organisations that develop, procure, deploy, or integrate AI systems.

Irish higher education institutions are implicated in the Act in two overlapping capacities. They act as deployers of AI systems procured from commercial vendors and integrated into teaching, assessment, learning analytics, and student support, and, in certain circumstances, as providers of AI systems developed internally and used beyond controlled research contexts. The Act also creates new rights for students, staff, and members of the public affected by AI-mediated decisions, which HEIs must respect in educational settings.

While the Act is technology-neutral in principle, its implications for teaching and learning are substantial. Educational contexts are explicitly recognised as sensitive domains where AI use can have lasting effects on opportunity, progression, and fairness. This places HEIs at the intersection of regulatory compliance, pedagogical design, and public accountability.

HEIs should recognise that alignment with the AI Act is not achieved only through classroom guidance or staff training. The systems most likely to create regulatory exposure in teaching and learning are procured, configured, and updated through vendor relationships. Research on the Act's educational and contractual implications stresses that institutions must treat transparency, informed consent, and liability as part of routine deployment governance, including what is specified in procurement documents and what is recorded as an institutional rationale for use. This matters for teaching because legal duties are shaped not only by what a tool can do in principle, but by how it is contractually embedded, operationally configured, and used in practice.¹

¹ Andrea Renda et al., *Study to Support an Impact Assessment of Regulatory Requirements for Artificial Intelligence in Europe: Final Report* (European Commission: Directorate-General for Communications Networks, Content and Technology, CEPS, ICF, Wavestone, 2021), <https://data.europa.eu/doi/10.2759/523404>; Sandra Wachter, 'Limitations and Loopholes in the EU AI Act and AI Liability Directives: What This Means for the European Union, the United States, and Beyond,' *Yale Journal of*



Recent legal scholarship on the EU AI Act stresses that its educational implications cannot be separated from its contractual ones.² AI deployment in education creates layered contractual relationships between institutions, students, educators, and third-party vendors, in which obligations relating to transparency, informed consent, and liability are central rather than ancillary. From this perspective, compliance is shaped by how responsibility is allocated when AI-mediated outputs affect assessment, progression, or educational opportunity. Uncritical reliance on AI tools risks eroding foundational knowledge and academic judgement, particularly in technically mediated disciplines, reinforcing the need for institution-level governance that treats procurement, disclosure, and pedagogical purpose as inseparable dimensions of lawful and responsible AI use in education.

This supporting instrument focuses specifically on teaching and learning. It does not attempt to provide a comprehensive institutional compliance manual, nor does it address research governance, commercialisation, or administrative automation except where these intersect directly with pedagogical practice. Its purpose is to assist HEIs in understanding how the Act applies to teaching and assessment, identifying areas of heightened regulatory exposure, and aligning pedagogical practice with both the letter and the spirit of the regulation.

Risk-based regulation and education

Rather than regulating all AI systems uniformly, the Act distinguishes between prohibited uses, high-risk uses subject to strict obligations, and lower-risk uses governed primarily through transparency and information duties.

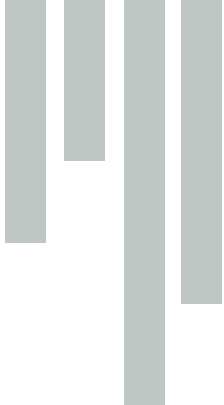
AI systems used in education are not automatically classified as high-risk. However, Annex III of the Act designates certain uses of AI in education and vocational training as high-risk where they have legal or similarly significant effects on individuals. These include systems used to:

1. determine access to or admission into educational programmes;
2. evaluate or grade learning outcomes where results materially affect progression or certification;
3. determine educational pathways or learning opportunities;

Law and Technology 26, no. 3 (2023),

<https://heinonline.org/HOL/Page?handle=hein.journals/yjolt26&id=671&div=15&collection=journals>.

² Ratko Ivković et al., 'The EU AI Act and Its Contractual and Educational Implications: AI, Education, and the Law of Obligations,' *International Journal of Cognitive Research in Science, Engineering and Education* 13, no. 2 (2025), <https://doi.org/10.23947/2334-8496-2025-13-2-551-562>.

- 
4. monitor behaviour during examinations or assessments;
 5. influence eligibility for educational support or opportunities.

The regulatory focus is therefore on function and effect, not on whether a system is labelled 'educational' or 'generative'. This distinction is critical for Irish HEIs when evaluating AI tools used across teaching and learning.

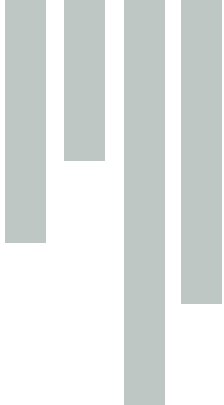
Most obligations applicable to high-risk systems take effect from August 2026, though certain prohibitions, including those relating to specific biometric and emotion recognition practices, apply earlier. Governance, enforcement, and sectoral guidance will continue to develop through secondary legislation and national regulatory coordination.

For Irish HEIs, compliance should be understood as an ongoing institutional process, not a single deadline-driven event. Teaching and assessment practices introduced or substantially modified from 2026 onward will need to be evaluated against the Act's requirements where they fall within scope.

AI systems in teaching and learning: practical implications

In higher education, the significance of the AI Act lies less in the presence of AI as such than in the educational functions to which it is applied. The Act does not treat educational technologies as inherently high-risk, nor does it regulate generative systems simply because they are used in teaching. Instead, it focuses on how AI systems operate within decision-making processes that carry material consequences for learners. For Irish HEIs, this functional orientation is crucial, as it requires institutions to look beyond labels and marketing claims and to examine what a system actually does in pedagogical practice.

Within teaching and learning, the uses most likely to trigger high-risk classification are those that evaluate, categorise, or constrain students in ways that shape progression or opportunity. Automated or semi-automated assessment systems that generate grades, marks, or evaluative judgements affecting academic standing sit squarely within this category. The same is true of AI-supported proctoring or invigilation tools that monitor student behaviour during examinations, particularly where these rely on biometric data or behavioural inference. Learning analytics systems that predict performance, identify 'at-risk' students, or recommend interventions can also fall within scope where their outputs influence access to support, progression routes, or educational opportunities. Decision-support systems used in admissions, pathway allocation, or recognition of prior learning raise similar concerns, as they intervene directly in processes that carry legal or similarly significant effects.



Much of the AI currently used in Irish higher education teaching operates at a lower level of regulatory risk. Generative tools that support formative feedback, drafting, practice exercises, or self-assessment typically do not, in themselves, determine outcomes or restrict opportunity. Writing assistants, coding tutors, and accessibility tools can often be used in pedagogically valuable ways without triggering high-risk obligations, provided that their role remains advisory, optional, and subject to human judgement. These systems are not unregulated, but the primary requirements concern transparency, appropriate disclosure, and alignment with existing data protection and quality assurance frameworks rather than the full compliance regime associated with high-risk applications.

Even where systems fall below the high-risk threshold, the regulatory direction of travel is towards demonstrable institutional competence rather than informal experimentation. Recent scholarship on the AI Act's implications for education treats transparency and human oversight as practical conditions of legitimate deployment, particularly in contexts where learners may over-trust fluent outputs or where staff may be nudged into automation bias when systems present recommendations with the veneer of objectivity. A proportionate institutional stance still requires clear student-facing disclosure, staff-facing guidance on system limits and failure modes, and a governance mechanism that records what an institution believes a tool is for, what it is not for, and what safeguards are expected in day-to-day use.³

The practical challenge for Irish HEIs is to distinguish clearly between uses that shape consequential decisions and those that support learning without determining outcomes. This distinction has direct implications for assessment design and the choice of educational technologies. Where AI systems are used in ways that affect grades, progression, or access, institutions must ensure that they meet the Act's expectations around documentation, oversight, and accountability. Where AI functions as a support for learning, institutions must still be transparent with students about its role and limitations, but can adopt a more proportionate governance approach.

This aspect of the Act aligns with a broader pedagogical direction already evident in Irish higher education. It encourages institutions to treat AI as a tool that can assist learning while reaffirming that academic judgement, assessment authority, and responsibility for educational outcomes remain human. The Act does not foreclose innovation in teaching and learning, but it does require that innovation be structured around clarity of purpose, careful differentiation of uses, and an explicit

³ Mirka Saarela et al., 'The EU AI Act: Implications for Ethical AI in Education,' in *Local Solutions for Global Challenges*, ed. Samir Chatterjee et al. (Springer Nature, 2025), https://doi.org/10.1007/978-3-031-93979-2_3.



understanding of where educational support ends and consequential decision-making begins.

Boundaries between research and teaching

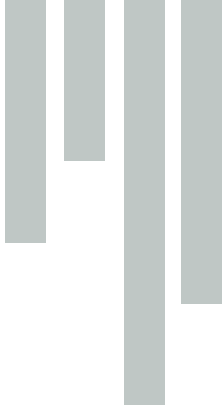
A recurring source of uncertainty for higher education institutions lies in the relationship between research activity and teaching practice under the AI Act. The regulation explicitly provides exemptions for AI systems developed and used exclusively for research and development purposes, reflecting a desire not to inhibit scientific inquiry or methodological experimentation. However, this exemption is tightly scoped and applies only while systems remain within controlled research contexts and are not placed into service in ways that affect individuals outside the research setting.

For Irish HEIs, this distinction is particularly salient because pedagogical innovation often emerges from research-led experimentation. Pilot projects, sandboxed trials, and exploratory uses of generative or adaptive AI may legitimately fall within the research exemption where participation is voluntary, impacts are limited, and outcomes do not affect grades, progression, or certification. The regulatory position shifts, however, once such systems are integrated into routine teaching, credit-bearing modules, or assessment practices. At that point, the system is no longer operating solely as a research artefact but as an educational technology with real consequences for students, and the protections afforded by the research exemption fall away.

Institutions need clear internal processes for identifying when an AI-enabled innovation moves from experimentation into teaching practice, and for reassessing its regulatory status at that point. Doing so allows innovation to continue while ensuring that deployment decisions are accompanied by appropriate governance and oversight once student outcomes are at stake.

Institutional responsibility and student rights

Under the AI Act, responsibility for AI use in teaching and learning is located primarily at institutional level, but it is realised through everyday pedagogical practice. Irish higher education institutions act as deployers of AI systems when they integrate commercial platforms, learning technologies, or internally developed tools into teaching, assessment, and student support. Individual educators do not become



regulated entities in their own right, but their choices and practices form part of the institution's compliance posture. This dual structure reshapes how pedagogical authority, oversight, and accountability are distributed across the sector.

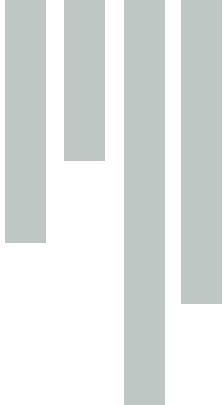
The Act presupposes that academic judgement remains central to teaching and assessment, even where AI systems are involved. Human oversight should be treated as an active condition of legitimate deployment. Where AI systems influence student outcomes, educators must be able to understand the purpose and limitations of those systems, interpret outputs critically, recognise risks such as automation bias, and intervene or override recommendations where pedagogically appropriate. In practice, this places constraints on how AI may be used in assessment workflows. Fully automated grading or evaluative decision-making without substantive human engagement is difficult to reconcile with the Act's expectations, particularly in high-stakes contexts affecting progression or certification.

These oversight obligations are inseparable from the rights that the Act confers on students. Where AI systems materially shape learning experiences, assessment outcomes, or progression decisions, students acquire enforceable entitlements to transparency and review. They must be informed when AI systems are in use, given intelligible information about how those systems function and are overseen, and able to request human review of AI-influenced decisions that affect them adversely. The Act does not require disclosure of proprietary technical details, but it does require that institutional practices be explainable in educational terms. Opacity undermines trust and, in regulated contexts, becomes a compliance risk.

For HEIs, this has direct implications for how teaching is communicated as well as how it is delivered. Course syllabi, assessment briefs, and learning analytics disclosures become regulatory artefacts as well as pedagogical ones. Students should be able to identify when AI systems are involved, what role they play, and where responsibility ultimately lies. This shifts transparency from a matter of institutional policy to a feature of course design and assessment communication.

The Act's treatment of emotion inference is particularly salient for learning platforms marketed as 'engagement', 'attention', or 'wellbeing' tools. Work on AI governance in the context of the EU AI Act highlights that education is a prominent application domain in current governance research, and that overreliance and misinterpretation are recurring risks when systems present behavioural or affective claims with a veneer of technical authority.⁴ In practical terms, HEIs should treat vendor claims about

⁴ Byeong-Je Kim et al., 'AI Governance in the Context of the EU AI Act,' *IEEE Access* 13 (2025), <https://doi.org/10.1109/ACCESS.2025.3598023>.



attention, affect, or emotional state as a red-flag category in teaching-and-learning procurement, and should require vendors to specify, in plain terms, whether a system attempts to infer emotion, how such inference is operationalised, and what safeguards or legal exceptions are relied upon.⁵

Certain practices, such as emotion recognition systems, are more sharply constrained, and tools that purport to infer attention, engagement, or affective states fall within this prohibition where they attempt to classify emotional or psychological characteristics. HEIs must review classroom analytics and engagement-monitoring technologies to ensure that prohibited practices are not embedded in teaching environments. More broadly, the Act signals a sceptical stance toward pedagogical surveillance, reinforcing the expectation that understanding student learning remains a matter of professional judgement rather than automated inference.

The Act does not diminish pedagogical autonomy, but it conditions that autonomy on transparency, oversight, and responsibility. For Irish higher education, this reinforces a model of teaching and learning in which AI may support academic work, but cannot displace the relational and interpretive dimensions of education that give academic judgement its legitimacy.

Assessment and learning analytics

The AI Act draws a sharp regulatory distinction between AI systems that support learning and those that participate in decisions with material consequences for students. In higher education, this distinction cuts across assessment, academic integrity, and student support, bringing these domains into a shared compliance frame. What matters is whether a system's outputs influence progression, opportunity, or institutional response.

In assessment, the Act reinforces a growing recognition that detection-led integrity strategies are both pedagogically weak and legally fragile. Automated tools that claim to identify AI-generated work or determine misconduct raise significant compliance concerns when they operate without transparent logic, meaningful oversight, or procedural safeguards. Systems that directly inform or automate disciplinary decisions are especially problematic. Decisions about academic integrity should remain

⁵ Wachter, 'Limitations and Loopholes in the EU AI Act and AI Liability Directives'; Saarela et al., 'The EU AI Act: Implications for Ethical AI in Education.'



academic decisions, embedded in disciplinary expertise and institutional procedure.

This regulatory logic extends naturally to assessment design. Tasks that rely exclusively on polished end products are increasingly vulnerable, both pedagogically and legally, in an AI-rich environment. By contrast, assessments that foreground process, explanation, and judgement align more readily with the Act's emphasis on accountability and oversight. Staged submissions, reflective components, oral explanation, and tasks that require students to articulate their reasoning or methodological choices make learning visible and preserve the centrality of human evaluation. In this sense, assessment redesign should be viewed as an essential part of a resilient response to a regulatory environment that prioritises explainability and responsibility.

Learning analytics and student support systems occupy a similarly sensitive position. Predictive models that identify students as 'at risk', recommend interventions, or prioritise support can offer genuine benefits, but they operate close to the threshold of high-risk classification when their outputs shape access to resources, progression pathways, or institutional attention. The key regulatory distinction again concerns whether systems are advisory or determinative. Analytics that prompt human review, inform reflective teaching practice, or support discretionary decision-making are more easily reconciled with the Act than systems that automate categorisation or trigger interventions without meaningful oversight.

This has practical consequences for how analytics infrastructures are designed and governed. Students must be able to understand how data about them are used and to seek review where AI-informed judgements affect their educational trajectory. Over-reliance on predictive models risks both regulatory non-compliance and pedagogical harm, particularly where such systems entrench bias or narrow conceptions of student success. The Act therefore encourages institutions to treat analytics as a support for professional judgement rather than as an instrument of automated decision-making.

The AI Act does not prohibit the use of AI in consequential educational contexts, but it raises the threshold of justification and oversight where outcomes matter. For Irish higher education, this convergence strengthens the case for pedagogical approaches that emphasise transparency and human responsibility. It also clarifies that the most robust responses to generative AI are those that redesign educational practices so that reasoning remains visible, even as AI tools are integrated into the fabric of teaching and learning.



Resource governance and capacity implications

Compliance with the AI Act carries material resource implications, particularly for high-risk applications. European Commission impact assessments suggest that regulatory obligations may impose non-trivial costs on organisations deploying AI systems, especially where governance structures, documentation practices, or technical expertise are underdeveloped. These figures are indicative rather than prescriptive, but they underscore that compliance is not cost-free.⁶

For Irish HEIs, investment will be required in staff development, procurement processes, documentation systems, and cross-functional governance linking teaching and learning, IT services, data protection, and legal oversight. These requirements are unlikely to be met through ad hoc or decentralised approaches alone. The Act creates incentives for collaboration. Shared services, coordinated procurement, and sector-level guidance can reduce duplication and allow institutions to pool expertise in areas such as vendor assessment, bias auditing, and compliance documentation. For a relatively small higher education system such as Ireland's, collective approaches offer a pragmatic route to meeting regulatory expectations without constraining pedagogical innovation. This becomes more pressing where governance research notes that regulated-system research remains relatively limited compared with the wider AI landscape, widening information asymmetries between what vendors claim and what deployers can evidence in practice.⁷

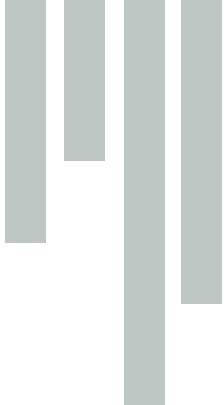
Proportionality and educational purpose

The AI Act does not prohibit the use of AI in higher education, nor does it mandate a single pedagogical response. It establishes conditions under which AI use must remain transparent, accountable, and subject to human judgement. Across the higher education sector, alignment with the Act should be understood as part of a broader commitment to public-interest education rather than as a narrow compliance exercise.

Universities occupy a dual role as both adopters of AI technologies and educators of future

⁶ Renda et al., *Study to Support an Impact Assessment of Regulatory Requirements for Artificial Intelligence in Europe: Final Report*.

⁷ Kim et al., 'AI Governance in the Context of the EU AI Act.'



professionals and citizens, and so how AI is integrated into teaching and learning carries normative weight. Modelling careful, proportionate, and values-led adoption is not only a regulatory requirement but a civic responsibility, particularly in a sector that shapes public understanding of knowledge and expertise.

As regulatory interpretation evolves, case law emerges, and pedagogical practice matures, HEIs will need to revisit their approaches to ensure continued alignment with both the formal requirements and the underlying aims of the AI Act. What the regulation ultimately demands is clarity of purpose, coherence of governance, and a renewed commitment to keeping learning, judgement, and responsibility firmly human.

References

- Ivković, Ratko, Dragan Vučković, Milan Ječmenić, et al. 'The EU AI Act and Its Contractual and Educational Implications: AI, Education, and the Law of Obligations.' *International Journal of Cognitive Research in Science, Engineering and Education* 13, no. 2 (2025). <https://doi.org/10.23947/2334-8496-2025-13-2-551-562>.
- Kim, Byeong-Je, Seunghoo Jeong, Bong-Kyung Cho, and Ji-Bum Chung. 'AI Governance in the Context of the EU AI Act.' *IEEE Access* 13 (2025). <https://doi.org/10.1109/ACCESS.2025.3598023>.
- Renda, Andrea, Rosanna Fanni, Moritz Laurer, et al. *Study to Support an Impact Assessment of Regulatory Requirements for Artificial Intelligence in Europe: Final Report*. European Commission: Directorate-General for Communications Networks, Content and Technology, CEPS, ICF, Wavestone, 2021. <https://data.europa.eu/doi/10.2759/523404>.
- Saarela, Mirka, Sachini Gunasekara, and Ayaz Karimov. 'The EU AI Act: Implications for Ethical AI in Education.' In *Local Solutions for Global Challenges*, edited by Samir Chatterjee, Jan vom Brocke, and Ricardo Anderson. Springer Nature, 2025. https://doi.org/10.1007/978-3-031-93979-2_3.
- Wachter, Sandra. 'Limitations and Loopholes in the EU AI Act and AI Liability Directives: What This Means for the European Union, the United States, and Beyond.' *Yale Journal of Law and Technology* 26, no. 3 (2023). <https://heinonline.org/HOL/Page?handle=hein.journals/yjolt26&id=671&div=15&collection=journals>.