

The Irish Journal of Technology Enhanced Learning

Volume 9 | Number 1

Article type: Original Research

Values-Led Governance of Generative AI in Irish Higher Education: An Interpretive Analysis of National Focus-Group Findings

James O'Sullivan

University College Cork & Higher Education Authority, Ireland james.osullivan@ucc.ie
<https://orcid.org/0000-0002-4214-9933>

Colin Lowry

Higher Education Authority, Ireland



This work is licensed under [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/)

Recommended Citation

O'Sullivan, James & Lowry, Colin (2026) Values-Led Governance of Generative AI in Irish Higher Education: An Interpretive Analysis of National Focus-Group Findings *Irish Journal of Technology Enhanced Learning*, 9(1). <https://doi.org/10.22554/366beh65>

Original Research

Values-Led Governance of Generative AI in Irish Higher Education: An Interpretive Analysis of National Focus-Group Findings

James O'Sullivan^{1,2*} & Colin Lowry¹

¹*University College Cork, Ireland*

²*Higher Education Authority, Ireland*

<https://doi.org/10.22554/366beh65>

Abstract

This paper examines how Irish higher education is attempting to govern generative AI at a moment when practice is evolving faster than policy. Drawing on an interpretive analysis of the Higher Education Authority's national consultation involving ten thematic focus groups and an institutional leadership summit held in early 2025, the study explores how staff, students, and sectoral leaders understand the pedagogical and ethical implications of large language models. The analysis re-reads the original coded transcripts through a governance lens, tracing how concerns about assessment, integrity, equity, and institutional readiness converge on deeper questions of educational purpose. Participants describe a sector marked by creativity and urgency but hampered by fragmented responses, uneven capacity, and uncertainty about how to translate shared values into coherent action. Attending to what is not said reveals further fault lines: environmental impact, labour conditions, research integrity, and commercial dependency remain largely absent from sectoral discourse despite their growing relevance. The paper argues that these omissions signal the need for a values-led governance framework capable of connecting policy, pedagogy, and infrastructure, and of sustaining reflexive, evidence-informed decision-making as technologies and practices change. Ireland's small and interconnected system is well placed to pursue such an approach, balancing national coordination with institutional autonomy. By foregrounding the values that shape sectoral perspectives — and the silences that accompany them — the paper identifies the conditions required for responsible, educationally coherent adoption of generative AI.

Keywords:

Generative AI, GenAI, policy, • Higher education • Educational policy • Digital pedagogy, Irish Higher Education

* Address for corresponding author james.osullivan@ucc.ie

The Irish Journal of Technology Enhanced Learning Ireland is the journal of the Irish Learning Technology Association, an Irish-based professional and scholarly society and membership organization. (CRO# 520231) <http://www.ilta.ie/>. This is an Open Access article distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0>).

1. Sectoral Perspectives on Generative AI

Since the public release of ChatGPT in November 2022, the presence of generative-AI tools within Irish higher education, particularly large language models (LLMs), within Irish higher education has become less a discrete technology than a pervasive condition, challenging assumptions about authorship, assessment, and by extension, the very purpose of higher education institutions. This transformation is unfolding with remarkable speed across Irish campuses, often outpacing institutional and national policy and governance, as well as pedagogical adaptation. Internationally, similar patterns of rapid and uneven adoption across higher education are being documented (Sutedjo et al., 2025, Šević et al., 2025, Jin et al., 2025).

In an effort to capture the perspectives of those most directly engaged in teaching and learning, the Higher Education Authority (HEA) convened a national consultation in early 2025 comprising ten thematic focus groups and an institutional leadership summit. These sessions brought together academic staff, students, professional-services personnel, and senior leaders from across the Irish higher-education sector, as well as representatives from public, semi-state, and industry stakeholders. The discussions offer a rare, ground-level view of how generative AI is being experienced in practice. By foregrounding the lived experiences of higher-education staff and leadership, the consultation surfaces the values, uncertainties, and priorities shaping Ireland's collective transition toward AI-mediated education. This sectoral consultation process was conducted as an evidence-gathering exercise specifically intended to inform the HEA's approach to generative AI in teaching and learning, and it fed directly into the development of the HEA's policy guidelines for the sector, *Generative AI in Higher Education Teaching & Learning: National Policy Framework* (O'Sullivan et al., 2025a), published in December 2025.

Key insights from the focus groups are summarised in *Generative AI in Higher Education Teaching & Learning: Sectoral Perspectives* (O'Sullivan et al., 2025). The analysis reveals that while the sector recognises that AI is reshaping education rapidly, institutional responses are perceived, consistent with international evidence, as fragmented and largely reactive. There is widespread acknowledgement of urgency, but at the time of the consultation, this had not yet translated into clear policy priorities or national strategic coordination. The resulting gap poses significant reputational and ethical risks that require attention at both institutional and national levels. Across all sessions, participants emphasise that generative AI fundamentally challenges core academic concepts including authorship, originality, assessment validity, and academic integrity. There is strong consensus that these principles require thoughtful re-articulation for an AI-mediated educational context, and the sector appears ready to engage with this complexity.

A particularly pressing concern centres on equity and inclusion. Without deliberate intervention, AI risks deepening existing educational divides (Smeaton et al., 2025). Unequal access to tools, inconsistent guidance, and varying levels of digital literacy threaten to widen inequities unless inclusion is designed into AI strategies from the outset rather than treated as an afterthought. Participants emphasise that inclusive practice requires structural commitment extending beyond the mere provision of tools.

Assessment emerges as a focal point of tension. Current dominant formats, particularly text-based take-home assignments, are increasingly vulnerable to generative-AI completion. Rather than advocating intensified surveillance or detection, participants express strong

appetite for redesigning assessment toward more authentic, process-focused, and dialogic forms. This represents a significant opportunity for pedagogical renewal, though institutional structures including rigid approval processes, workload pressures, and risk aversion pose substantial barriers to innovation.

Related concerns about academic integrity reveal that traditional definitions of misconduct are inadequate for addressing AI-generated work. The sector grapples with how to distinguish between reasonable and inappropriate use of large language models such as ChatGPT. Over-reliance on unreliable detection tools raises serious questions about due process, while students report widespread confusion about what constitutes appropriate use. Participants emphasise that integrity cannot be instilled through rules and detection alone but must be embedded through course design, classroom dialogue, and staff modelling. There is strong support for moving toward a partnership model with students, treating them as co-creators of norms rather than subjects of policing.

Significant anxiety surrounds the potential erosion of foundational academic skills, particularly writing and critical thinking. Staff observe that students increasingly use generative AI to complete tasks rather than learn through them, particularly writing and critical thinking. Writing is repeatedly framed as a cognitive process essential for organising thought and developing disciplinary understanding, and generative AI use risks reducing it to a hollow deliverable, severing it from its pedagogical function.

The consultation identifies a critical gap in AI literacy across both student and staff populations. Participants emphasise that effective AI literacy extends beyond tool proficiency to include critical evaluation, understanding of limitations and biases, ethical reasoning, and informed judgement about when not to use such tools. Students already use generative AI extensively but often uncritically, without appreciating implications for their learning. Staff, particularly those outside technical disciplines, feel under-prepared to teach or model AI literacy and lack confidence in their own understanding. This points to what participants identify as perhaps the most significant limiting factor in responsible AI adoption, staff capacity and cultural readiness. Academic staff are expected to make rapid pedagogical shifts with minimal training, unclear policy guidance, and insufficient time. Without substantial investment in professional development, communities of practice, and protected time for adaptation, sustainable change is unlikely.

In terms of infrastructure, participants describe a patchwork of provision across institutions, with some holding exploratory licences, others relying on free tools, and many operating without guidance. This fragmentation creates inconsistency, increases data-privacy risks, and reinforces digital inequity. Participants stress that infrastructure encompasses far more than technical platforms, extending to governance processes, procurement ethics, support services, shared practices, and the organisational capacity to use tools appropriately. Governance structures are described as fragmented and reactive, and often lacking credibility. The absence of clear institutional frameworks forces ethical decision-making downward onto individual staff members, creating anxiety and inconsistency. Students perceive double standards, while staff express exhaustion and confusion. Leaders acknowledge that institutions face multiple vectors of risk, from unethical grading and diminished graduate competence to erosion of public and employer trust. There is strong consensus that governance must prioritise educational values alongside technical and legal compliance, requiring cross-functional collaboration among IT, academic development, student services, and teaching staff.

Leaders recognise that generative AI presents sophisticated governance, reputational, and cultural challenges that require institutions to revisit their core purpose and values. Participants urge institutions to lead proactively, embedding critical thinking, ethical literacy, and social responsibility within generative AI strategies. Leadership, they argue, must articulate what kind of learning and education matter in an AI-enabled world.

The consultation reveals strong appetite for coordinated national action, provided such coordination balances sectoral coherence with institutional flexibility. Participants advocate shared frameworks, pooled resources, and clear principles rather than rigid, prescriptive policies that cannot adapt to rapid technological change or diverse institutional contexts. Examples from other jurisdictions where universities collaborate on standards, toolkits, and infrastructure are cited as potential models. While a range of cross-institutional initiatives and collaborations exist within the Irish system, the analysis here does not seek to catalogue or evaluate specific programmes; instead, it focuses on how consultation participants articulated the need for coordination and what forms of coordination were imagined as legitimate and workable.

While each of the focus group sessions focused on a specific theme, several strategic priorities consistently emerge across each session:

1. Reframe the challenge around educational purpose and values rather than treating generative AI primarily as a compliance or security issue;
2. Support assessment and curriculum innovation through funding, exemplars, and legitimised experimentation;
3. Understand staff and student capacity-building as pedagogical and ethical work, not merely technical training;
4. Design inclusion into AI strategies from the outset, ensuring marginalised voices shape implementation;
5. Strengthen governance structures to enable transparent, values-led decision-making;
6. Develop infrastructure holistically, encompassing people, processes, and principles alongside platforms.

Throughout the consultation, participants emphasise the importance of engaging students as partners rather than problems. Students want honest conversations about generative AI, clear guidance on appropriate use, and opportunities to contribute to norm-setting. This partnership approach is seen as essential for ethical and practical reasons, as it builds trust and improves adherence to integrity standards, while better preparing students for a world already transformed by generative AI.

The report concludes that Irish higher education stands at a critical juncture. The sector demonstrates readiness for deliberate, values-led engagement with generative AI and possesses considerable creative capacity for adaptation. Realising this potential requires moving beyond fragmented, reactive responses toward coordinated national frameworks that provide clarity while respecting institutional autonomy and diversity. Success depends on sustained investment in staff capacity, genuine partnership with students, values-based governance, and leadership willing to articulate educational purpose in an age of intelligent technologies. The shift from detection to design, from compliance to partnership, and from management to stewardship represents the sector's preferred direction of travel, one that positions Ireland as an emerging leader in ethical, pedagogically sound AI integration.

This ‘direction of travel’ also aligns with a wider shift in how jurisdictions and public institutions are framing generative AI governance, away from narrow compliance responses and towards principles-based, values-led governance paired with risk-oriented organisational strategy. At the policy level, influential international instruments emphasise trustworthy AI grounded in human rights and democratic values, including human agency and oversight, transparency and explainability, accountability, and robustness and safety (OECD, 2019; UNESCO, 2021).

In parallel, operational governance approaches increasingly stress the translation of these principles into institutional controls and decision criteria, such as documented risk assessment, context-sensitive use-case differentiation, evaluation and monitoring, and the integration of AI governance into established organisational processes including teaching and assessment design, data governance, procurement, and quality assurance (Tabassi, 2023). Within Europe, this trajectory is further reinforced by the adoption of a risk-based regulatory framework for AI that foregrounds proportionality and oversight, while heightening the salience of system-level concerns (European Parliament & Council of the European Union, 2024). Read in this context, the Irish consultation’s emphasis on values, coordination, and stewardship can be situated within a recognisable international governance paradigm, one that also provides a defensible basis for examining which issues did, and did not, consolidate as cross-cutting priorities in the consultation discourse despite prominence in contemporary AI governance frameworks (UNESCO, 2021).

2. Methodology

This paper undertakes a secondary, interpretive analysis of the focus-group findings published in *Generative AI in Higher Education Teaching & Learning: Sectoral Perspectives* (O’Sullivan et al., 2025b). The original consultation was designed and facilitated by the HEA to capture and summarise cross-sectoral perspectives on the implications of generative AI in Irish higher education. The consultation comprised ten thematic focus groups delivered online using a standardised facilitation format, alongside a final in-person institutional leadership consultation, each structured around semi-structured prompts aligned to a predefined thematic domain.

This current study draws on the summary analysis, as well as directly on the verbatim transcripts of those sessions, which were anonymised and coded to preserve confidentiality while retaining the authenticity of participant voice. Participants in the original consultation were recruited via targeted invitations through institutional contacts and sectoral networks, with an intention to capture diversity of role—including academic staff, students and student representatives, professional services, and institutional leaders—institution type, and discipline; the consultation report indicates that the participant cohort included early adopters and cautious observers and that all institutions under the HEA’s remit were represented.

This analysis adopts an interpretive approach, re-reading the original coded data through the lens of policy provision, situating the descriptive findings of the original report within a broader epistemic and governance framework, identifying where sectoral perceptions align with, or diverge from, emerging global evidence and discourse. The analytic starting-point for this re-reading is the consultation’s established thematic structure and cross-cutting findings, which function in this paper as the ‘present themes’ against which absences are interpretively identified, or, what was repeatedly foregrounded across sessions, versus what was weakly articulated, compartmentalised, or treated as peripheral.

This paper identifies interpretive tensions and thematic absences that speak to deeper value conflicts within the Irish higher-education response to generative AI. ‘Absent themes’ are operationalised as issues that are prominent within contemporary policy and governance discourse on generative AI, nationally or internationally, but which did not appear as sustained, cross-cutting concerns in the consultation transcripts, or appeared only episodically within bounded segments without shaping participants’ dominant problem framings or policy demands. The interpretive orientation of this paper recognises that meaning is co-constructed between participants, analysts, and the wider discursive environment in which generative AI is being negotiated. The focus-group transcripts are therefore treated not only as reflections of practice but as artefacts of a moment in cultural and institutional transition. The aim is to generate insight into how values are contested and operationalised across the sector rather than to measure prevalence or consensus.

To strengthen transparency, the secondary analysis followed a structured interpretive workflow. A renewed familiarisation pass was undertaken across the full set of anonymised transcripts and the original summary analysis, with analytic memos developed to capture stated priorities and rationales, recurring tensions or contradictions across roles and sessions, and preliminary candidates for ‘silences’, that is, issues expected to surface given the governance stakes of generative AI. The original coding structure applied to the consultation transcripts was then re-examined to distinguish issues that were coded as central, issues that were coded but marginal, and issues that were not meaningfully present, noting where discussion clustered around immediate operational concerns—such as assessment vulnerability, integrity enforcement, AI literacy, infrastructure—rather than longer-horizon governance considerations. Candidate absent themes were subsequently tested against the transcripts through a structured ‘absence check’, using targeted re-search and negative-case probing to identify counterinstances and clarify the conditions under which a theme did surface. The resulting pattern of emphases and silences was synthesised as interpretive tensions and thematic absences, and situated within a wider governance and epistemic frame to clarify implications for values-led national strategy. On this basis, the secondary analysis is interpretive rather than confirmatory: it does not seek to audit the original coding, but to treat it as a stable evidential base for a focused re-reading oriented to policy and governance meaning.

This analysis acknowledges its situatedness. The authors are engaged in national and institutional work on AI policy and educational development in Ireland, and this positionality informs the interpretation. In addition to acknowledging positionality, an explicitly reflexive stance is adopted: proximity to national policy development may sensitise analysis to particular governance framings and may also risk over-weighting policy-salient themes when interpreting stakeholder discourse. To mitigate this, a transcript-first analytic sequence is retained before introducing external governance frames; memoing and negative-case searching are used to avoid over-claiming absence where bounded discussion occurred; and ‘absence’ is treated as a matter of cross-cutting salience rather than literal non-occurrence. The intention is to contribute to collective sectoral understanding by interpreting stakeholder perspectives within an evolving international and ethical context, highlighting the governance challenges that demand coordinated, values-led responses, and then acting upon those interpretations through the development of appropriate frameworks.

3. Absent Themes

While the national focus groups surface immediate pedagogical and governance concerns, a critical reading reveals several significant thematic absences that warrant attention. Such absences may reflect consultation design and the stated scope of inquiry, and what emerges in qualitative research is shaped by what is invited, what feels relevant to participants in the moment, and what falls within the perceived boundaries of discussion. Absent topics may therefore be as much products of the research encounter as they are indicators of gaps in sectoral thinking, and they merit interpretation as such.

Nonetheless, these omissions highlight potential blind spots in how Irish higher education currently frames its response to generative AI. Patterns of omission often point to areas where values remain unexamined, or where structural conditions render certain issues invisible. In policy and governance analysis, attending to silence can therefore expose the limits of a sector's imagination, what it struggles to resource or confront, and where guidance and recommendations issued by public bodies can have the most impact.

Environmental sustainability

Environmental sustainability is absent from the thematic focus-group transcripts analysed in this study. Outside the concluding leadership summit—where it featured as part of the programme design rather than as a participant-led topic—it did not emerge within the focus-group discussions. Despite growing evidence that training and operating large language models carry substantial carbon costs and consume significant water and energy resources (Bender et al., 2021, Luccioni et al., 2023), environmental impact receives only passing mention. This silence is particularly jarring given that Irish higher-education institutions have made public commitments to climate action and sustainable development. The failure to foreground ecological consequences as a core ethical consideration represents a troubling disconnect between institutional sustainability rhetoric and the framing of generative AI as primarily a pedagogical and integrity challenge.

There is no discussion of whether institutions might favour providers with stronger sustainability credentials, whether energy consumption should feature in procurement criteria, or whether AI literacy for students and staff should include awareness of the environmental costs of digital technologies. The omission is striking given that environmental sustainability is one of the few areas where ethical principle and practical constraint naturally align. Unlike domains in which values and convenience pull in opposite directions, the environmental case for moderation in AI adoption could reinforce more thoughtful and pedagogically grounded implementation, but this framing is absent, suggesting that sustainability is not integrated into participants' sense of what responsible generative AI adoption entails.

Labour and wellbeing

The consultation documents widespread staff anxiety and exhaustion but stops short of analysing these as structural or ethical concerns requiring institutional responsibility. References to staff feeling 'underprepared' or 'burnt out' appear largely as implementation challenges rather than as symptoms of deeper issues around workload intensification and the affective toll of rapid pedagogical change. It is telling that during focus-group recruitment

many staff expressed interest in contributing but were unable to participate because of unsustainable workloads. Missing from the findings is any sustained engagement with an ethics of care perspective, one that reframes staff and student wellbeing as foundational to responsible AI adoption rather than as a secondary consideration.

This omission extends to questions of labour and professional identity. While participants acknowledge that graduates will enter ‘a workforce transformed by AI’, there is little critical reflection on what this transformation implies for the nature and security of academic work itself. The possibility that generative AI might be mobilised to rationalise reduced contact hours, larger class sizes, or further casualisation of teaching roles goes largely unexamined. Most of the focus group sessions treat AI as a challenge to pedagogy and integrity but not as a significant threat to labour conditions within the sector.

Mental health and wellbeing, though mentioned in the inclusionary-practices session, receive only passing attention. The documented links between academic pressure and student mental health are not explored in depth, nor is the psychological burden created when institutional expectations around generative AI use remain unclear or contradictory. These gaps suggest that wellbeing continues to be framed as an individual concern rather than a collective ethical responsibility integral to a values-led approach.

Research and epistemic autonomy

While the HEA’s policy development has focused explicitly on teaching and learning, it remains striking how little attention is given to research and to the domains where research and teaching explicitly intersect. The binary between the two is, in any case, largely artificial, since all or certainly most higher-education teaching should be research-led. The consultation does treat them as distinct spheres, reflecting the HEA’s strategic decision to focus first on teaching and learning, where the immediate sense of urgency and disruption is most acute across the sector. Still, little is said about research supervision, the ethics of AI-assisted literature reviews, the use of generative tools in data analysis, or the integrity challenges facing postgraduate researchers.

Equally absent is philosophical reflection on what generative AI means for conceptions of knowledge and the purpose of education itself. Discussions remain largely pragmatic, centred on ‘how to respond’ rather than ‘what is at stake’. Few participants consider how generative AI might alter the relationship between knower and known, or whether certain forms of cognitive labour such as synthesis and critique are under particular threat.

Related to this is the question of academic freedom and epistemic autonomy. Several participants voice concern that governance responses might over-regulate AI use without sufficiently protecting scholarly independence or interpretive diversity. Risk-averse policy can easily stifle legitimate experimentation or impose top-down uniformity on disciplines that depend on contextual judgement.

The focus groups surface a tension between the need for coordinated governance and the preservation of academic autonomy but leave unresolved how higher education might safeguard its epistemic integrity in an era when authorship itself is increasingly mediated by machines.

Second-degree plagiarism

Of particular concern to the HEA's policy development is the problem of second-degree plagiarism, the reproduction or summarisation of AI-generated material that itself draws, often invisibly, on unattributed human sources (O'Sullivan & Lowry, 2025). Such recursive reuse corrodes the principle of intellectual traceability on which academic authorship rests, blurring the distinction between synthesis and appropriation. This issue goes without mention across the focus groups, and left unchecked, it risks normalising derivative knowledge production that is fundamentally at odds with the university's commitment to originality, transparency, and critical inquiry.

Governance and commercial power

While participants express strong support for national frameworks and inter-institutional coordination, the consultation produces few concrete proposals for how such collaboration might be funded or governed. The gap between aspiration, such as the call for shared infrastructure, and implementation remains unresolved. Questions of which body should lead coordination, how institutional diversity can be respected within a national approach, and what mechanisms of accountability or enforcement might apply are largely deferred.

This absence is compounded by limited critical engagement with the commercial dynamics shaping generative AI adoption in higher education. Although procurement is mentioned, the consultation does not interrogate the growing influence of technology vendors in defining educational priorities, nor the risks of vendor lock-in and data dependency. The possibility that educational practice might be reconfigured to fit commercial product architectures rather than pedagogical need receives little scrutiny. Issues such as data extraction, opaque terms of service, and the monetisation of learning interactions are acknowledged but not typically represented as threats to institutional autonomy or the public interest.

Learning from experience and building evidence

The consultation makes no reference to failed experiments or cautionary examples from other jurisdictions. While international frameworks are frequently cited as models, there is little reflection on what has not worked elsewhere, or on the importance of creating space for Irish institutions to experiment, fail, and learn without reputational penalty. This omission reflects a wider cultural pattern within higher education, the stigmatisation of failure and the accompanying tendency to adopt innovation uncritically and defensively rather than as a process of iterative inquiry.

No mechanism is proposed for continuous evaluation or longitudinal monitoring of generative AI's educational or social impacts. The consultation generates a rich snapshot of a moment in transition, but without infrastructure for ongoing evidence-building, there is a danger that policy will rest on static assumptions as technologies and practice, and accompanying risks, evolve rapidly. Key questions, such as who should conduct such monitoring, how data might be gathered ethically and comparably across institutions, and what thresholds would trigger review or revision, remain unanswered. A genuinely values-led approach institutionalises reflexivity, embedding evaluation and learning as enduring features of governance.

4. Conclusions & Policy Implications

The findings of the HEA's national consultation make clear that the sector perceives generative AI as a force that is reconfiguring the conditions under which knowledge is produced, and consequently, the conditions under which it must be assessed and governed. Across the focus groups, participants demonstrate awareness of this transformation and a strong appetite for collective guidance, an appetite that is now reflected in the HEA's *Generative AI in Higher Education Teaching & Learning: National Policy Framework* (O'Sullivan et al., 2025a). This work does not exist in a policy vacuum and sits alongside, and must be implemented through, the HEA's established policy commitments and system objectives, including sustainability, equity, and wellbeing across the higher-education system. The governance challenge is therefore as much one of integration and coherence, aligning generative AI guidance with existing policy architecture and institutional practice, as it is one of generating new policy statements. The discussions also reveal fragmentation and uncertainty, and institutional responses have often remained reactive and uneven, shaped by local capacity rather than coordinated vision. This imbalance reflects a wider international pattern in which higher-education systems recognise the magnitude of AI-driven change but have yet to define its governance as an ethical and strategic responsibility rather than a technical one.

Generative AI forces universities to revisit the question of what constitutes learning and originality, disrupting inherited distinctions between process and product, and challenges assessment systems built on evidence of individual cognition. This pressure to redesign assessment is increasingly reflected in the higher-education literature, which argues that generative-AI capability accelerates the shift from product-only tasks toward more authentic, process-evidencing, and context-specific assessment designs that can better evidence student learning under conditions of ubiquitous tool use (Xia et al., 2024; Miao & Holmes, 2023). The focus groups show that Irish educators are already grappling with these questions in classrooms and committees, often without adequate institutional support. The risk is that decisions about pedagogy and integrity will continue to be made in an ad hoc manner, leaving values implicit and inconsistently applied. In particular, the documented limitation of AI-writing detection tools, including false positives and vulnerability to evasion, make detection-led enforcement a poor foundation for academic-misconduct processes, raising acute fairness and due-process risks if used as decisive evidence rather than as weak, contextual signals (Weber-Wulff et al., 2023; Giray, 2024). International guidance cautions that routine substitution of generative tools for core academic practices can weaken the development of foundational skills—especially writing as a cognitive process, disciplinary judgement, and critical engagement—unless curricula and assessment are deliberately designed to preserve and evidence those capabilities (Miao & Holmes, 2023). The opportunity, however, lies in recognising generative AI as a catalyst for renewal, prompting a deeper re-articulation of higher education's public purpose, and in treating the HEA's national policy framework as an enabling baseline for coherent institutional implementation rather than as an endpoint.

What is missing from the focus groups underscores this point: sustainability, labour, research integrity, and commercial power remain under-examined precisely because they sit at the intersection of ethics and governance. Their omission signals the need for a broader understanding of gen AI adoption as a social and institutional process shaped by competing priorities, limited resources, and inherited hierarchies of value. Any values-led governance framework must account for how generative AI is used and, perhaps more importantly, for

what is rendered invisible in its implementation. This involves expanding the policy imagination beyond compliance and capability toward stewardship, where universities act as custodians of both knowledge and public trust, and ensuring that implementation remains attentive to domains that are explicitly recognised within the HEA's framework principles, but which can be deprioritised in day-to-day operational response.

Such a framework requires embedding ethical reasoning, reflexivity, and evidence-building into the governance cycle itself. Evaluation must be treated as a continuous, participatory process through which policy adapts to technological and cultural change. Sustainability metrics, workload equity, and epistemic autonomy should feature alongside efficiency and innovation. This reorientation would also demand that governance structures model the very literacies they seek to cultivate, most notably, transparency and accountability, consistent with the framework's emphasis on sector learning, monitoring, and iterative revision in response to emerging evidence and experience.

For Ireland, this represents both a challenge and an opportunity. The country's relatively small and interconnected higher-education system is well positioned to pursue coordinated, values-led experimentation. A national approach can provide coherence without uniformity, enabling institutions to develop contextually appropriate practices within shared ethical parameters. The HEA, working with representative bodies and agencies, is well placed to convene this process and to translate sectoral dialogue into actionable frameworks, and has now provided a national reference point intended to be adapted locally rather than enforced as a uniform rule set. Internationally, there is a clear opening for Ireland to distinguish itself by integrating generative-AI policy with commitments to equity, sustainability, and democratic governance. Any values-led national strategy should therefore:

1. Integrate AI literacy and critical pedagogy across disciplines so that students and staff can engage with the technology reflectively and responsibly;
2. Invest in sustainable and secure digital infrastructure supported by transparent procurement and accountable data governance;
3. Embed students and staff as partners in co-designing curricula, assessment, and ethical norms;
4. Align institutional policies with adaptive, evidence-informed regulatory principles capable of evolving with technological and social change;
5. Monitor environmental, labour, and epistemic indicators to ensure that AI adoption advances socially just and educationally coherent outcomes.

The governance of generative AI in higher education is not about managing risk or enforcing compliance but about shaping the moral and intellectual direction of digital transformation. A values-led approach reframes policy as a form of collective interpretation, through which the sector continually negotiates what it means to learn and teach in the digital age. If Ireland can act deliberately on this recognition, anchoring its strategy in reflection and care, it will help define how a democratic education system can thrive within generative AI contexts, with the HEA's national framework providing the shared orientation through which that deliberation can be coordinated and sustained.

The findings of this analysis point to the need for a coordinated national approach to generative AI that is explicitly grounded in educational values and social responsibility. Such an approach should treat AI not as an operational problem to be managed but as a governance challenge requiring alignment across policy, pedagogy, and infrastructure, and should be

evaluated primarily through the quality and consistency of institutional implementation and review, rather than the existence of policy statements alone.

References

- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)* (pp. 610–623). Association for Computing Machinery.
<https://doi.org/10.1145/3442188.3445922>
- European Parliament, & Council of the European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828. Official Journal of the European Union, OJ L, 2024/1689.* <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- Giray, L. (2024). The problem with false positives: AI detection unfairly accuses scholars of AI plagiarism. *The Serials Librarian, 85*(5–6), 181–189.
<https://doi.org/10.1080/0361526X.2024.2433256>
- Jin, Y., Yan, L., Echeverria, V., Gašević, D., & Martinez-Maldonado, R. (2025). Generative AI in higher education: A global perspective of institutional adoption policies and guidelines. *Computers and Education: Artificial Intelligence, 8*, 100348.
<https://doi.org/10.1016/j.caeai.2024.100348>
- Luccioni, A. S., Viguier, S., & Ligozat, A.-L. (2023). Estimating the carbon footprint of BLOOM, a 176B parameter language model. *Journal of Machine Learning Research, 24*(253), 1–15. <https://jmlr.org/papers/v24/23-0069.html>
- Miao, F., & Holmes, W. (2023). *Guidance for generative AI in education and research*. UNESCO. <https://doi.org/10.54675/EWZM9535>
- OECD. (2019). *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449). Organisation for Economic Co-operation and Development.
<https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449>
- O’Sullivan, J., & Lowry, C. (2025). *Ten considerations for generative artificial intelligence adoption in Irish higher education*. Higher Education Authority.
<https://doi.org/10.5281/zenodo.14917857>
- O’Sullivan, J., Lowry, C., Woods, R., & Conlon, T. (2025a). *Generative AI in higher education teaching and learning: National policy framework*. Higher Education Authority. <https://doi.org/10.82110/px37-mp48>
- O’Sullivan, J., Lowry, C., Woods, R., Marrinan, B., & Hutchinson, C. (2025b). *Generative AI in higher education teaching and learning: Sectoral perspectives*. Higher Education Authority. <https://doi.org/10.5281/zenodo.17153423>

- Popović Šević, N., Šević, A., Slijepčević, M., & Krstić, J. (2025). AI adoption in higher education: Exploring attitudes and perceived benefits between users and non-users. *Online Journal of Communication and Media Technologies*, 15(4), e202528. <https://doi.org/10.30935/ojcm/17246>
- Smeaton, A., Ahern, D., Birhane, A., Leavy, S., Riordan, B., & O'Sullivan, B. (2025, February). *AI and education* (AI Advisory Council Advice Paper). Government of Ireland, Department of Enterprise, Tourism and Employment. <https://www.gov.ie/en/department-of-enterprise-tourism-and-employment/publications/ai-advisory-council-advice-papers/>
- Sutedjo, A., Liu, S. P., & Chowdhury, M. (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>
- Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
- UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. United Nations Educational, Scientific and Cultural Organization. <https://unesdoc.unesco.org/ark:/48223/pf0000380455>
- Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P., & Waddington, L. (2023). *Testing of detection tools for AI-generated text*. *International Journal for Educational Integrity*, 19, 26. <https://doi.org/10.1007/s40979-023-00146-z>
- Xia, Q., Weng, X., Ouyang, F., Lin, T. J., & Chiu, T. K. F. (2024). A scoping review on how generative artificial intelligence transforms assessment in higher education. *International Journal of Educational Technology in Higher Education*, 21, 40. <https://doi.org/10.1186/s41239-024-00468-z>