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

Generative AI in Higher Education Teaching and Learning

Sectoral Perspectives

Contributors

James O'Sullivan
Colin Lowry
Ross Woods
Brian Marrinan
Clyde Hutchinson

Thank you to our colleagues across the Irish higher education sector who took part in this research.

How to cite:
O'Sullivan, James, Lowry, Colin, Woods, Ross, Marrinan, Brian, and Hutchinson, Clyde. *Generative AI in Higher Education Teaching and Learning: Sectoral Perspectives*. Higher Education Authority. 2025. https://www.teachingandlearning.ie/gen-ai-sectoral-perspectives/

This report is licensed under the Creative Commons Attributions 4.0 International license (CC BY 4.0). Please provide attribution and link to the Higher Education Authority as stated above.

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

Foreword

This report presents the findings of a national consultation on generative AI in Irish higher education, commissioned by the Higher Education Authority (HEA) as part of a values-led, evidence-informed approach to policy development in this rapidly evolving area. The consultation represents a critical, time-sensitive step in developing Ireland's national policy guidelines for generative AI in higher education teaching and learning.

While individual institutions are developing local responses, the sector currently lacks the coordinated national policies needed to ensure consistency, equity, and educational integrity across the system, a gap this consultation will help address. The HEA's forthcoming policies and guidelines on generative AI in higher education will be informed by international frameworks and sectoral scholarship, but it is also essential to draw from local, discipline-specific perspectives and experiences, which was the major function of this consultation process. This ensures that Ireland's response to generative AI is grounded in global best practice while remaining responsive to national educational values and priorities.

As part of its broader commitment to sectoral support, the HEA, through its National Forum for the Enhancement of Teaching and Learning, has developed a dedicated Artificial Intelligence in Irish Higher Education resource portal, which offers centralised access to national and international policies, institutional guidelines, and open educational resources focused on teaching and learning in AI-mediated contexts.

The HEA has deliberately chosen a values-led approach to policy development—as evidenced by its *Ten Considerations for Generative Artificial Intelligence Adoption in Irish Higher Education*—recognising that generative AI raises fundamental questions about the purpose, process, and integrity of higher education in Ireland. Rather than pursuing a purely technical or compliance-driven response, this approach centres values like academic integrity, critical thinking, inclusion, sustainability, and human agency as the foundation for policy development.

The perspectives gathered through the ten thematic focus groups and institutional leadership summit documented in this report provide essential insight into how generative AI is currently being experienced across Irish higher education. Participants—including academic staff, students and student representatives, professional support personnel, and institutional leaders—shared their practices, concerns, aspirations, and recommendations for navigating this transformative moment.

Their contributions will directly inform forthcoming policy guidelines, ensuring that HEA recommendations are grounded in the lived realities of those teaching, learning, and leading within the sector. The conversations documented in this report demonstrate the sector's readiness to engage creatively and conscientiously with the opportunities and challenges generative AI presents, while also highlighting the urgent need for coordinated support and clear guidance.

This consultation marks an important step toward developing a coherent national response to generative AI in higher education. The insights gathered reveal both the complexity of the challenge and the sector's capacity for thoughtful adaptation. As the HEA and higher education institutions build upon these findings through continued collaboration and policy development, this report serves as both a snapshot of the present moment and a foundation for shaping Ireland's educational future in the age of generative AI.

James O'Sullivan
Teaching & Learning Policy Advisor for Artificial Intelligence
Higher Education Authority

Colin Lowry
Senior Manager for Teaching and Learning, Enhancement, and Digital Innovation
Higher Education Authority

1. Table of Contents

2. Executive Summary

Generative Artificial Intelligence (AI) is already reshaping the landscape of higher education. Irish institutions are actively engaging with its implications for teaching, learning, assessment, academic integrity, inclusion, and governance. Yet responses to date have largely been emergent and institution specific, with many educators and students adapting in real time, and institutions facing difficult questions about what constitutes appropriate, ethical, and effective use.

This report presents the findings of a national consultation commissioned by the Higher Education Authority (HEA), involving ten thematic focus groups and a leadership summit. The sessions brought together academic staff, students, professional support personnel, and institutional leaders to explore their experiences, concerns, and aspirations regarding the role of generative AI in higher education.

The consultation surfaced a strong appetite for structured dialogue, sectoral guidance, and capacity-building. While participants recognised both the risks and opportunities of generative AI, several consistent themes emerged across all sessions:

1. Strategic Coordination Is Urgently Needed

   The sector is moving rapidly but unevenly. To prevent capacity fragmentation, a coordinated national approach is needed.

2. Educational Purpose Should Be Re-examined

   AI challenges fundamental values of authorship, assessment, and academic integrity. There is an opportunity for re-articulation of these values in the context of generative AI, but not abandonment.

3. Inclusion Requires Deliberate Action

   AI can widen existing inequalities unless strategies explicitly centre accessibility, equity, and Universal Design for Learning (UDL).

4. Assessment Reform Is Both Necessary and Welcome

   Participants called for a shift away from detection as a primary strategy and towards authentic, process-based assessment supported by AI-literate practices.

5. Staff And Students Need Structured Support

   Professional development, communities of practice, and student partnerships will be essential to sustainable, ethical adoption.

6. Governance And Infrastructure Should Be Holistic

   Beyond technical tools, institutions require clear ethical frameworks, procurement guidance, and cross-functional leadership on AI adoption.

7. Leadership Has a Role in Shaping, Not Just Managing

   The AI transition is one which will require strategic direction and national coherence to build capacity and deliver successfully on opportunities.

The report identifies a set of strategic considerations to inform national policy development and institutional planning. These include:

- Developing values-based practice frameworks aligned with the core values and goals of Irish higher education
- Supporting programme-level curriculum and assessment redesign
- Investing in staff capacity, student partnership, and inclusive digital literacy
- Strengthening governance structures to manage AI adoption transparently and ethically
- Balancing institutional diversity with system-wide coherence through shared infrastructure and resources

This report offers a grounded, system-wide picture of how generative AI is currently being experienced across Irish higher education. It surfaces the complex realities of change already underway, and provides a set of reflective insights for policymakers and institutional leaders as they navigate a rapidly evolving educational environment.

AI is not just a technological issue, it is a pedagogical, ethical, and strategic one. This consultation has made clear that Ireland's higher education sector is ready to move forward with care, clarity, and creativity. Going forward, a shared vision, structured support, and coordinated response will support sustaining educational values while embracing innovation.

3. Introduction and Background

3.1 Project Context and Rationale

Generative artificial intelligence (GenAI) is a subfield of artificial intelligence that uses generative models to produce text, images, videos, or other forms of data, and its rapid development has the potential to reshape the global landscape of higher education. With the widespread public availability of tools such as such as Open AI's ChatGPT and other emerging generative AI platforms across disciplines, the boundaries of authorship, learning, assessment, and academic integrity are being redefined in real time. These developments are not occurring in the margins as they are now central to how students learn, how staff teach, and how institutions uphold educational purpose.

Recognising the magnitude of this shift, the Higher Education Authority (HEA) initiated a national consultation to explore how generative AI is perceived, used, and understood across Irish higher education. This project sought to move beyond technological hype and instead create space for informed, critical, and inclusive dialogue.

The timing of this work reflects a pivotal juncture. Institutions are being called upon to respond to AI with clarity and leadership, yet many lack shared frameworks, coordinated infrastructure, or sufficient capacity-building support to respond consistently. Equally, students and staff are already navigating AI in their daily practices, yet institutional guidance often remains unclear or still in development. Without a deliberate and collective response, the sector risks fragmenting along lines of institutional capacity, staff confidence, and student access.

This report contributes to the development of an informed national strategy by bringing forward the voices of those most directly engaged in teaching, learning, governance, and student support. It aims to inform the HEA and sectoral stakeholders as they shape coordinated, values-led responses to the ongoing integration of AI into higher education.

3.2 Research Objectives

The project was commissioned to explore the lived realities, concerns, aspirations, and strategic priorities of key stakeholders in relation to generative AI. The objectives were to:

- Understand how AI is currently being used by students, academic staff, professional staff, and institutionally across teaching, learning, assessment, and student engagement.

- Identify risks and challenges that stakeholders associate with AI, including issues of academic integrity, ethics, inclusion, skills development, and infrastructural gaps.
- Surface opportunities and innovations where AI is seen to support access, feedback, writing development, teaching enhancement, or inclusive learning.
- Explore institutional readiness in terms of governance, culture, leadership, and capacity for staff and student development.
- Generate insight to inform national policy, funding strategies, and future guidelines on the responsible and effective integration of AI in higher education.

3.3 Key Questions Framing the Research

The focus groups were guided by a set of framing questions adapted to each session's topic and participant group. These included exploring how generative AI is currently being used in educational settings, what the perceived benefits, risks, and ethical dilemmas associated with its use are, the ways AI is changing teaching, learning, writing, assessment, the forms of guidance or support that are currently in place, and where there are gaps.

The inquiry was qualitative by design, aiming not to produce consensus or fixed recommendations, but to surface the complexity, nuance, and lived expertise that underpin responsible decision-making.

4. Methodology

This section sets out the approach used to design, conduct, and analyse a national programme of focus groups commissioned by the HEA to explore the evolving role of generative AI within Irish higher education. The research was undertaken to provide a credible, nuanced, evidence base to inform national policy, institutional guidance, funding priorities, and sector-wide coordination.

The methodological design prioritised both analytical rigour and practical relevance. Particular emphasis was placed on capturing the lived realities of academic staff, students, professional support staff, and institutional leaders who are currently encountering generative AI in varied and often experimental ways. This approach reflects the HEA's strategic need to understand, not only high-level trends, but also emerging ground-level practices, concerns, and opportunities.

4.1 Purpose and Rationale

The study was designed to provide a grounded, detailed understanding of how generative AI is currently perceived, used, and debated across Irish higher education institutions. Given the pace of technological change and the absence of settled norms, the research sought to explore a wide spectrum of attitudes, ranging from cautious experimentation to active adoption, alongside the structural, ethical, and pedagogical challenges arising from AI integration.

Rather than treating generative AI as a single, monolithic development, the study was structured to investigate how different aspects of AI intersect with core academic and institutional functions. The aim was not only to document current experiences but to surface the assumptions, values, and tensions that may shape future trajectories in teaching, learning, governance, and infrastructure.

Qualitative methods were chosen to enable depth, complexity, and reflexivity. Focus groups, in particular, were selected for their capacity to foster peer-to-peer dialogue, encourage the articulation of implicit assumptions, and generate insight through contrast and debate among participants with different roles and perspectives.

4.2 Focus Group Structure and Thematic Design

A total of eleven focus groups were conducted. Ten sessions were delivered online using a standardised format and facilitation protocol. Each of these lasted approximately two hours. A final, in-person session was held which focused specifically on strategic and governance-level issues from the perspective of institutional leaders.

Each session explored a distinct thematic domain relevant to generative AI in higher education. This thematic structure enabled a deeper investigation into specific questions while still allowing for the emergence of cross-cutting insights.

HEA policy advisors designed priority themes after a period of sectoral engagement. The themes of the sessions were as follows:

1. Enhancement for Teaching – Investigating how generative AI is used to streamline, support, or enhance the teaching process across planning, delivery, and feedback.
2. Enhancement for Assessment – Exploring how AI is influencing assessment design, feedback generation, summative evaluation, and perceptions of validity.
3. Skills and Cognitive Impact – Examining whether and how generative AI might erode or transform core academic and graduate-level skills.
4. AI Literacy – Considering how AI literacy might be defined, developed, and supported for both staff and students.
5. Ethical AI – Surfacing concerns about fairness, bias, and the need for institutional safeguards and ethical governance.
6. Academic Integrity – Focusing on evolving definitions of originality, misconduct, and the boundaries of acceptable AI support.
7. Writing and Authorship – Debating whether AI-assisted writing changes the pedagogical function of writing, and how norms might need to evolve.
8. Inclusion Practices – Assessing whether generative AI can help personalise support and reduce barriers, or whether it risks reinforcing exclusion.
9. Infrastructure Readiness – Exploring the capacity of institutions to adapt to AI from technical, financial, policy, and cultural standpoints.
10. Student Experience, Teaching & Learning – A sector-level conversation with the SETL Committee to explore the impact specifically on this topic.
11. Strategic Leadership and Governance – The research culminated in a leadership consultation, part of the HEA Leadership Summit on Generative AI in Higher

Education Teaching & Learning, hosted at University College Cork, addressing system-wide risks, coordination needs, and long-term responsibilities.

Each focus group followed a consistent format: an initial orientation, a series of guided discussion prompts aligned to the theme, a "practice exchange" segment in which participants shared specific tools or experiments they had tried or observed, and a closing reflection.

4.3 Participant Composition and Recruitment

Across the eleven sessions, more than 80 participants contributed, representing academic staff, students, professional services, and external stakeholders from across the Irish higher education sector (see Appendix 2). All HEIs under the HEA's remit were represented, alongside non-profits, student organisations, and industry partners. While the group was diverse, the majority of participants identified as lecturers and/or researchers. This emphasis is appropriate: lecturers occupy the frontline of teaching and assessment, the areas most immediately reshaped by generative AI. Their pedagogical decisions determine how institutional policies are enacted in practice, how academic integrity is safeguarded, and how disciplinary standards are maintained. For this reason, ensuring lecturer perspectives were well represented was essential to producing findings that are both realistic and implementable.

Participants were recruited through targeted invitations issued via institutional contacts and sectoral networks. Care was taken to ensure diversity of institution type (e.g. universities, technological universities, institutes of technology), discipline, role, and level of experience with generative AI. Participants were invited to disclose their background and level of familiarity with generative AI as part of an introductory warm-up activity, helping the facilitation team tailor the discussion accordingly. Generally, the participants identified as being early adopters or cautious observers with an average to above average grasp of GenAI compared to their peer group.

The inclusion of mixed-role groups was intentional as this study sought to capture multi-perspective conversations that reflect the realities of collaboration, policy development, and implementation within higher education institutions. The final in-person session brought together senior leaders from institutions, providing a strategic overview and sectoral perspective on the themes identified across the earlier sessions.

4.4 Data Collection and Ethical Considerations

All focus groups were facilitated by experienced qualitative researchers using a consistent semi-structured discussion guide tailored to the session's theme.

Participants were provided with information sheets and gave informed consent in advance of each session. Sessions were recorded (audio only) with permission, and full transcripts were produced for analysis. All data were handled in accordance with GDPR and anonymised before analysis to protect participant confidentiality.

Discussion prompts were designed to be open-ended and exploratory, allowing participants to share their experiences, speculate on future developments, and interrogate the implications of different forms of AI use. Special attention was paid to fostering an environment of mutual respect, curiosity, and critical reflection, particularly given the fast-moving and occasionally polarising nature of the topic. Participants did not always agree, divergent perspectives were observed and highlighted where they arose.

A distinctive feature of the data collection design was the "practice exchange" segment. In each session, participants were invited to share examples of how they or their colleagues had used AI tools in real or experimental ways, ranging from generating quiz questions or simplifying explanations to using large language models to write policy drafts or personalise student feedback. These use cases provided concrete insight into the current frontier of experimentation and helped ground theoretical discussion in lived practice.

4.5 Analytical Framework, Process, and Limitations

The data were analysed using a thematic analysis approach, following the established framework of Braun and Clarke (2006[1]), a method for identifying and reporting patterns (themes) within qualitative data. This method was selected for its flexibility and capacity to support both descriptive and interpretive insights. The process included the following stages:

1. Data Familiarisation: Each transcript was read in full, and reflective memos were created to capture initial patterns, contradictions, and points of interest.
2. Coding: A coding framework was developed through a combination of inductive (emerging from the data) and deductive (aligned with HEA's policy focus) approaches. Transcripts were coded to ensure systematic tracking of ideas, patterns, and language.
3. Theme Generation: Thematic clusters were developed from related codes and refined through iterative review. Where appropriate, sub-themes were identified to capture nuance.

[1] Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*

4. Cross-Case Synthesis: Matrices were used to examine how themes varied across roles, institutions, and thematic sessions. Particular attention was given to identifying tensions, contradictions, and areas of silence.
5. Interpretation and Validation: Findings were interpreted in light of wider sectoral developments and strategic policy goals. Draft analyses were reviewed internally for coherence, representativeness, and clarity. Peer debriefing and inter-coder reliability checks were used to enhance rigour and minimise bias.

The use of direct participant quotations throughout the report serves to illuminate findings and ensure the authenticity of participant voices. These have been lightly edited for clarity and anonymised to preserve confidentiality.

As a qualitative study, findings are not statistically generalisable to the entire sector. However, the diversity of roles, disciplines, and institutions represented, coupled with the deliberate thematic structure of the sessions, provides a high level of analytic generalisability and practical relevance. The fast-moving nature of the generative AI landscape means that the findings reflect a snapshot in time, albeit one that captures both the early contours of institutional adaptation and the enduring values and dilemmas likely to shape future policy.

5. Cross-Cutting Themes and Thematic Findings

This section presents the principal findings of the national consultation on generative AI in Irish higher education. It draws on ten thematic focus groups and a dedicated forum with institutional leadership.

The findings are presented in two parts:

- **Part 1: Cross-Cutting Themes**

  This synthesis distils a set of shared challenges, insights, and priorities that emerged consistently across multiple sessions. These themes highlight the systemic nature of generative AI's impact on higher education, extending across pedagogy, assessment, infrastructure, inclusion, institutional readiness, and public trust. They offer a strategic overview of the issues facing the sector and form the foundation for the national-level considerations included later in the report.

- **Part 2: Individual Session Findings**

  Thematic summaries of each session follow, including focused discussions on topics such as teaching enhancement, academic integrity, AI literacy, assessment design, and inclusive practice. Each summary outlines the key concerns, opportunities, and recommendations arising within that particular domain, and reflects the diversity of perspectives contributed by participants.

Together, these findings offer a rich, grounded, and forward-looking evidence base to support the Higher Education Authority and Irish higher education institutions in shaping a coordinated, values-led response to the integration of generative AI.

5.1 Cross-Cutting Themes Across All Sessions

In the table below we outline eight themes that were common across several of the groups.

Theme Title	Cross-Cutting Theme Description
Urgency Without Prioritisation	*"We can't eat the whole elephant, so let's pick off a section and really focus on that."* Across staff, student-facing professionals, and institutional leaders, there is a shared recognition that generative AI is reshaping higher education at speed, yet current responses are often reactive, fragmented, and inconsistent. Participants stressed the need to define clear priorities, focus on manageable areas, and coordinate strategically to avoid reputational, educational, and ethical risks.
Redefinition of Core Concepts	*"We can't keep pretending that old definitions still work. But we can't abandon them either."* AI is blurring boundaries around authorship, originality, assessment, learning processes, and academic integrity. Many participants felt that core values such as e.g. academic integrity, intellectual honesty, and equity, remain essential, but could be re-articulated in ways that make sense in an AI-mediated world.
Equity and Inclusion Are at Risk Without Intervention	*"If you don't centre inclusion, you're building new walls, not opening doors."* While AI can be viewed as a promising tool for accessibility, language support, and learning scaffolds, participants highlighted growing concern that unequal access to tools, guidance, and digital literacy risks deepening divides. Inclusion cannot be assumed, rather, it should be designed.
Assessment Is a Flashpoint, But Also a Leverage Point	*"We can't rely on detecting AI use, we need to change how we assess, full stop."* Participants across all roles recognised that AI renders many assessment formats vulnerable. But this was also viewed as an opportunity to rethink what we assess, how, and why. There is

	widespread consensus for more authentic, process-focused, and AI-literate forms of assessment.
Staff Feel Underprepared and Under-supported	*"There's an assumption that staff can just absorb this. But people are exhausted, and this is a pedagogical shift, not just a tech one."* Participants are expected to redesign practice in real time, but lack policy clarity, professional development, and time. This is not simply a technical gap; it is a structural and cultural readiness issue that could enable effective, ethical change.
Students Want to Be Engaged Partners	*"They want to talk about this. They don't want to be policed, they want to be prepared."* There is strong support for open dialogue with students as partners and co-creators in developing AI guidelines, rather than treating AI use only as a disciplinary matter. AI use is already widespread among students, and they want honest, ethical conversations about what good learning should now look like, ultimately leading to better adherence to integrity standards.
Infrastructure Is More Than Tools	*"We can't build this on free trials and workarounds."* Licensing, platforms, and IT capacity matter, but so do governance models, procurement ethics, support services, and shared practices. There is a balancing act with holistic, integrated, and values-led infrastructure, as well as cost, privacy, and security.
Leadership Has a Role in Shaping, Not Just Managing	*"Leadership needs to be brave enough to say what education is for, not just what policies are needed."* From the summit and session 10, a strong message emerged: leadership should shape, not merely manage, the AI transition. This means aligning educational purpose, institutional values, and national coherence, and building the cultural capacity to act with clarity and care.

5.2 Individual Session Findings

5.2.1 Focus Area: Enhancement for Teaching

The focus group on enhancement for teaching revealed both optimism and uncertainty around the use of generative AI in pedagogy. Participants viewed AI as a promising tool for improving teaching preparation, content generation, and personalisation. However, they also expressed serious concerns about staff deskilling, the uncritical adoption of AI outputs, and a growing disconnect between AI experimentation and pedagogical intention.

Several participants emphasised that meaningful integration of AI into teaching requires more than technical skill. It demands deliberate educational design, reflective practice, and a clear institutional framework that balances innovation with academic values. Without these, staff risk using AI in ways that are expedient rather than educationally sound.

Key Insights	AI is already being used to support common teaching tasks, especially in large, content-heavy modules.
	Staff are unsure how to balance efficiency with pedagogical integrity when using AI-generated materials.
	There is a notable absence of institutional guidance, examples, or dialogue about quality and responsible use.
	AI-enhanced teaching remains an individual effort, often isolated from curriculum development processes or student input.

Implications for Policy and Strategy	Institutional and national frameworks should explicitly address the pedagogical dimension of AI, not just policy or compliance.
	Investment is needed in communities of practice that support reflective and intentional teaching innovation with AI.
	Consideration of AI would benefit from integration into existing learning design, curriculum review, and professional development pathways.
	Policy should clarify the boundary between AI assistance and the human role in teaching, clarifying the position on the outsourcing of pedagogical judgment.

5.2.2 Focus Area: Enhancement for Assessment

The assessment-focused session surfaced deep concern across disciplines about the relevance and resilience of current assessment practices in the age of generative AI. Participants described traditional assignments, particularly text-based, take-home work, as increasingly vulnerable to AI-generated completion. Yet rather than resort to enforcement or technological detection, most participants expressed a clear preference for pedagogical and design-based responses.

Staff articulated the need to reimagine assessment to prioritise process, reflection, student voice, and original thinking. AI was seen not as a threat to academic standards, but as a stimulus to re-centre learning. However, structural constraints, limited capacity, and fear of institutional risk aversion were seen as major barriers to change.

Key Insights	AI exposes the fragility of existing assessment formats, and also highlights their limitations.
	There is strong appetite to redesign assessment towards more authentic, formative, and dialogic forms.
	Institutions have not yet consistently provided the leadership or support needed to enable widespread innovation.
	Assessment change is entangled with broader issues of workload, governance, and trust.

Implications for Policy and Strategy	The shift from detection and compliance toward pedagogical transformation is an important transition for assessment policy.
	National agencies can play a leadership role in supporting assessment redesign through guidance, exemplars, and funding.
	Investment may be needed in assessment literacy, resourcing innovation, and normalising experimentation in assessment practice.

5.2.3 Focus Area: Foundational Skills and Skills Degradation

This session explored concerns about the long-term impact of generative AI on students' development of foundational academic and cognitive skills — including writing, critical thinking, information literacy, and independent learning. Participants acknowledged that while AI can scaffold or support some learning tasks, its widespread and uncritical use may erode the very capabilities higher education is intended to cultivate.

The group emphasised that foundational skills are not merely functional competencies but form the basis of disciplinary identity, intellectual autonomy, and lifelong learning. Several participants expressed concern that students might bypass key developmental experiences, arriving at outputs without engaging with underlying processes. The potential for 'skill decay' was discussed not only in students, but in staff themselves.

Key Insights	Generative AI poses a credible risk of eroding students' foundational writing, thinking, and reasoning capacities if used passively or excessively.
	There is a lack of clarity around which skills will still be taught, practiced, and assessed in an AI-rich environment.
	Staff fear students may become over-reliant on AI tools for surface performance, without deeper understanding.
	Institutional messaging rarely distinguishes between learning with AI and learning via AI substitution.

Implications for Policy and Strategy	Clear, discipline-informed statements are needed on which foundational skills remain essential for higher education, and how AI should support (not replace) their development.
	Assessment strategies should explicitly engage with skill development processes, not just outputs.
	Institutions should design AI guidance and teaching practices that build student awareness, reflection, and self-regulation around tool use.

5.2.4 Focus Area: AI Literacy

Participants in this session stressed that developing AI literacy is now essential for both students and staff, not only to use generative AI tools effectively, but to question them critically. AI literacy was broadly defined to include not just technical proficiency, but also ethical awareness, critical thinking, understanding of AI's capabilities and limitations, and informed judgment about when and how to use such tools.

Despite wide agreement on the urgency of developing these literacies, there was confusion about who should lead this work, how it should be embedded in curricula, and what minimum standards or outcomes should look like. Participants were concerned that without a coherent strategy, AI literacy risks being optional, uneven, or reduced to simplistic tool use.

Key Insights	An effective approach to AI literacy treats it as a shared responsibility across the institution, rather than leaving it to isolated champions or optional workshops.
	There is a need to go beyond 'how to use the tool' and include critical evaluation, social context, and ethical reasoning.
	Students are already using AI tools, often uncritically, and risk developing dependency without deeper understanding.
	Staff feel underprepared to teach or model AI literacy, particularly in disciplines outside computer science.

Implications for Policy and Strategy	National and institutional frameworks should consider the definition of AI literacy holistically and whether to promote it as a core competency for graduates and educators alike.
	Curriculum strategies need to consider embedding AI literacy within subject learning, rather than treating it as a standalone or optional topic.
	Institutions would benefit from staff training, co-creating learning materials, and involving students in shaping literacies for an AI-mediated world.

5.2.5 Focus Area: Ethical AI Use in Higher Education

This session revealed a strong sense of ethical ambiguity, institutional lag, and concern about accountability in the use of generative AI. Participants recognised the significant power of AI tools but questioned whether their use aligns with higher education's core values, including fairness, transparency, intellectual integrity, and inclusion.

Participants called for a shift away from solely individual responsibility (student or staff) towards shared, institutionally supported ethical frameworks. The lack of clarity on boundaries, what is permissible, what is encouraged, and what is exploitative, creates a vacuum filled by guesswork, risk aversion, or silent adoption. There was clear appetite for institutions to lead by articulating principles, not just rules.

Key Insights	Staff are navigating ethical grey areas without sufficient guidance or shared language.
	Students may use AI tools in ways that reinforce inequity or circumvent learning, but not always maliciously or even understanding the impact.
	There is concern about the ethics of the tools themselves, including how they are trained, monetised, and governed.
	A values-led, participatory approach is needed to shape ethical norms and behaviours.
Implications for Policy and Strategy	Ethical use of AI in higher education can be addressed through principles, dialogue, and collective responsibility, not compliance alone.
	Institutions need to convene diverse voices to co-create ethical standards and scenarios for using generative AI in learning and teaching.
	Policy responses benefit from including staff, student, and societal perspectives. AI ethics is not just a technical or procedural matter.

5.2.6 Focus Area: Academic Integrity in the Age of Generative AI

Participants expressed concern that generative AI is reshaping understandings of academic integrity in ways that outpace institutional responses. While traditional definitions of plagiarism and misconduct still apply in some cases, AI introduces new grey zones where student intent, process, and disclosure are not easily observable. There is, therefore, doubt around enforcement ability, with stories of legal cases and other challenges abounding.

There was strong consensus that integrity should not be reduced to detection or enforcement. Instead, participants called for a renewed focus on academic values, clarity of expectations, and partnership with students in shaping new norms. Importantly, staff emphasised that fostering a culture of integrity in an AI-enabled world requires transparency, trust, and the alignment of policy, pedagogy, and assessment design.

Key Insights	Existing academic integrity frameworks are not well-equipped to address generative AI use.
	Detection-focused approaches are insufficient on their own if used without due process or transparency.
	Students need clear, educative guidance on what constitutes acceptable AI use, and why.
	Integrity should be seen as a shared cultural value, not just a set of rules to be enforced.

Implications for Policy and Strategy	Institutions should shift from a punitive to a pedagogical model of integrity that aligns with assessment design and student development.
	Academic integrity policies require updating to address generative AI explicitly, including clear definitions, disclosure practices, and examples.
	National leadership is needed to support harmonisation across institutions and reduce inconsistency in expectations.

5.2.7 Focus Area: Writing in an AI-Rich Educational Environment

Participants described writing as one of the areas most directly and visibly impacted by generative AI tools. While acknowledging potential benefits such as scaffolding, accessibility, and productivity, the session revealed deeper tensions about the purpose, process, and pedagogy of writing in higher education.

There was concern that AI use could undermine students' development of writing as a cognitive and disciplinary practice. Participants also questioned the sustainability of current teaching and assessment models built around writing outputs, particularly when AI can now generate text that appears fluent but lacks original thought. Combining oral assessments, supervised in-class work, and take-home work provides an interesting, if resource intensive, approach to assessment, as well as requiring structured, screen-free reading and research experiences. The group called for a reframing of writing as a process, something to be taught, supported, and made visible, rather than a final product to be evaluated.

Key Insights	Writing remains central to disciplinary learning, identity, and assessment, but its purpose is now contested.
	AI is being used by students for everything from idea generation to full draft creation, often invisibly.
	Staff lack clarity on what to permit, discourage, or assess — and feel existing guidance does not capture the nuance.
	There is a need to redesign curricula and pedagogy around writing as a developmental, reflective, and critical act.
Implications for Policy and Strategy	Institutions should reaffirm the role of writing as a learning process, not merely a deliverable.
	Guidance is needed to help staff and students navigate AI use in writing, distinguishing between support, substitution, and misconduct.
	Writing development should be more fully integrated into programme design, with space for reflection, iteration, and dialogue about AI.

5.2.8 Focus Area: Inclusionary Practices and Generative AI

This session surfaced a complex, and sometimes contradictory, set of views about the relationship between generative AI and inclusion in higher education. On the one hand, AI tools were seen as potentially empowering, offering support for students with disabilities, multilingual learners, or those with gaps in prior academic preparation. On the other hand, participants warned that AI could reinforce exclusion and inequality when access, skills, or cultural assumptions are not taken into account.

The central insight was that AI does not automatically promote or undermine inclusion, rather, its impact depends on how institutions embed it into pedagogy, policy, and student support. Without deliberate strategies, there is a real risk that generative AI will widen existing disparities and deepen disengagement among marginalised learners.

Key Insights	AI offers new possibilities for inclusive practice but also introduces new forms of exclusion if not critically implemented.
	Students with disabilities, neurodivergence, or language needs may benefit from AI tools, but only with clear, supported, and ethical use.
	There are growing gaps in digital literacy, confidence, and access across the student body, and few institutional mechanisms to address them.
	Inclusion ought to be proactively designed into AI strategies, not treated as an assumed benefit.

Implications for Policy and Strategy	In designing AI-related pedagogy, assessment, and infrastructure, institutions benefit from placing inclusion at the centre.
	Student support services, accessibility units, and digital learning teams should be resourced and integrated into AI planning.
	Inclusion strategies should address structural inequality, not only individual adaptation.

5.2.9 Focus Area: Infrastructure for Generative AI in Higher Education

Participants identified infrastructure as either a foundational enabler, or constraint, in realising the educational potential of generative AI. Discussions went beyond hardware and software to include digital governance, institutional capabilities, procurement ethics, data protection, and cross-functional resourcing. The group highlighted significant uncertainty, inconsistency, and under-preparedness at both institutional and national levels.

Participants emphasised that without coherent, scalable, and equitable infrastructure, AI adoption would be piecemeal, risk-laden, and likely to deepen divides between institutions and learners. Staff voiced the need for clear leadership on licensing, access, and interoperability, as well as strategic alignment between digital learning, IT, library, and teaching teams.

Key Insights	Current infrastructure is fragmented, reactive, and often bypassed through individual or unofficial tool use.
	There is a widespread lack of institutional guidance on what tools are approved, safe, or pedagogically aligned.
	Ethical concerns (e.g. student data, licensing, and commercial influence) are compounded by weak procurement oversight.
	Infrastructure will include not only tools but the human, organisational, and policy systems that support their use.

Implications for Policy and Strategy	There is a question as to whether AI infrastructure should be viewed as a public academic asset, requiring coordinated, transparent, and values-based management.
	Institutions need national-level support and guidance to ensure security, equity, and pedagogical coherence.
	Without strategic investment, AI risks becoming a peripheral or shadow system outside institutional control.

5.2.10 Student Engagement and Teaching & Learning in the Age of Generative AI

This session, which included representatives with a strong institutional remit for student engagement and teaching enhancement, brought a systems-level lens to the challenges and opportunities of generative AI. Participants reflected on the urgent need to adapt teaching and learning models, not only to accommodate AI, but to re-centre pedagogy, student motivation, and inclusive design in a changing landscape.

While participants acknowledged that AI tools offer powerful supports for both learners and educators, there was concern that education systems are reacting rather than leading, with limited coordination, capacity, or vision. A clear message emerged in that AI is not an external disruption to be managed, but a structural force that invites redesign of how higher education approaches knowledge, learning, and engagement.

Key Insights	AI challenges not only the delivery of content but the fundamentals of how learning is framed, valued, and assessed.
	There is an urgent need for pedagogical leadership and staff development that extends beyond basic digital training.
	Students can be engaged not simply as users or risks, but as co-creators of learning in an AI-mediated world.
	Without structural and cultural shifts, institutions risk superficial adoption, gradual erosion of ethical standards or oversight, and student disengagement.

Implications for Policy and Strategy	AI should be positioned within broader strategies for pedagogical innovation, curriculum renewal, and student engagement.
	Institutional policies that are enabling, rather than restrictive, should be encouraged, along with supporting experimentation, interdisciplinary approaches, and inclusive practice.
	National-level coordination is needed to provide direction, resources, and shared infrastructure for transformative change.

5.2.11 Focus Area: Institutional Leadership Perspectives on Generative AI

The leadership summit surfaced a strategic tension between urgency and uncertainty. Institutional leaders recognised that generative AI is a systemic force reshaping not only pedagogy, but the purpose, values, and operation of higher education. However, most institutions lack a coherent roadmap, shared language, or internal alignment to respond adequately.

Leaders expressed concern over fragmented governance, growing risks to reputation and trust, and the widening gap between national ambition and local capability. Participants called for a sector-wide approach that integrates infrastructure, values, staff capacity, and student partnership, while avoiding reactive or technocratic solutions. Above all, there was consensus that leadership can shift from a compliance mindset to one of educational stewardship.

Key Insights	AI challenges core institutional values including integrity, autonomy, and the role of human learning.
	There is a need for collective infrastructure, sectoral alignment, and nationally supported guidelines.
	Leadership should act decisively, ethically, and inclusively, even amid uncertainty, to avoid erosion of trust, credibility, and educational value.
	Institutional readiness is not only technical; it includes cultural, ethical, academic, and communicative dimensions.

Implications for Policy and Strategy	National coordination is faced with balancing consistency with local flexibility, developing a shared framework, not a single policy.
	Institutions should invest in continuous staff and student education around AI, ethics, and changing practices.
	Governance should prioritise not simply risk containment, but clarity, integrity, and critical engagement.

6. Implications and Considerations for Policy and Practice

The ten thematic focus groups and institutional leadership summit revealed a sector actively grappling with generative AI's implications, but still facing substantial uncertainty and disparate levels of readiness and capacity. While this report does not offer prescriptive recommendations, a number of key considerations arise from the shared insights and dilemmas expressed across the consultation. These are offered as prompts for future policy development and institutional planning efforts led by the HEA and the wider sector. These considerations will inform the forthcoming national policy framework on AI; they outline key considerations and discussions for policymakers and institutions.

Reframing Educational Purpose in the Age of AI

There is growing recognition that generative AI is not merely a tool to be integrated, but a force that invites higher education to revisit its underlying educational values. Across sessions, participants called for institutions and policy bodies to reflect on what remains essential in human learning, and how this is to be protected, redefined, or reimagined.

Considerations:

- What values should shape how AI is used in teaching and learning, and how can these be made explicit?
- How might the purposes of assessment, academic integrity, and student development be reframed to reflect both continuity and change?
- In what ways can national and institutional strategies support critical, reflective, and ethical engagement with AI, rather than compliance-driven implementation?

Supporting Pedagogical and Assessment Innovation

Assessment was the most frequently cited site of tension, experimentation, and urgency. Participants saw AI as exposing the limitations of legacy assessment practices, but also creating new openings for process-focused, authentic, and inclusive approaches.

Considerations:

- How can the sector be supported to design forms of assessment that are not only resilient to AI, but enriched by it?
- What types of funding, time, or peer support might enable academic staff to trial and evaluate new practices?
- What role might national guidance or sector-wide exemplars play in legitimising innovation while respecting disciplinary autonomy?

Investing in Staff and Student Capacity

While awareness of AI is widespread, confidence in using it pedagogically, ethically, or critically remains uneven. Staff consistently expressed a need for support, space, and shared reflection. Students, too, voiced a desire for clearer expectations, more open dialogue, and structured opportunities to learn with and about AI.

Considerations:

- How can professional development be framed to build pedagogical, not just technical, capacity around AI?
- What structures (e.g. communities of practice, interdisciplinary learning circles) might help normalise ongoing adaptation?
- In what ways can students be included as partners in shaping AI literacy, norms, and educational expectations?

Embedding Inclusion and Accessibility in AI Strategy

Participants emphasised that AI adoption is not neutral, it may exacerbate existing inequities unless actively counterbalanced. The need to centre inclusion, accessibility, and Universal Design principles were raised across discussions on writing, teaching, assessment, and infrastructure.

Considerations:

- What processes are needed to ensure that AI-related initiatives do not unintentionally disadvantage certain student groups?
- How can institutions ensure that marginalised or underrepresented voices shape AI-related decision-making and practice?

- What shared resources or practices could support inclusive implementation across a diverse sector?

Strengthening Institutional Governance and Infrastructure

Infrastructure was defined broadly, encompassing platforms and tools, yes, but also governance processes, procurement decisions, and the cultural readiness of institutions. Many reported a lack of clarity about who is responsible for leading on AI internally, and how decisions are being made.

Considerations:

- How can institutions ensure that governance of AI reflects educational values as well as technical or legal compliance?
- How can the sector maintain public trust in the integrity of qualifications and learning outcomes in an age of AI-assisted student work?
- What cross-functional structures (e.g. ethics boards, curriculum steering groups) might help coordinate decisions?
- To what extent could national frameworks help establish shared baselines without constraining local responsiveness?

Promoting Sectoral Coherence While Respecting Institutional Diversity

There was strong appetite for coordinated action, shared principles, pooled resources, national frameworks, but also recognition that institutions differ in size, mission, and capacity. Many participants called for a flexible model of alignment, where national bodies support coherence without prescribing uniformity.

Considerations:

- What form might a shared national framework for AI in higher education take, and how can it balance adaptability with clarity?
- Where might it be appropriate to pool infrastructure, guidance, or training across the sector?
- How can local experimentation be surfaced, shared, and supported to inform broader system learning?

Anticipating Future Developments While Responding to Present Needs

Finally, participants across all roles acknowledged that the pace of change is unlikely to slow. Many warned against premature policymaking or rigid controls, but equally cautioned that failure to act could undermine credibility and public trust.

Considerations:

- How can national and institutional strategies remain agile, revisable, and responsive to future developments?
- What foresight or scenario-planning mechanisms might help the sector prepare for what lies ahead?
- How can trust be built and maintained among students, staff, and the wider public, as AI becomes more deeply embedded in academic life?

This section is not intended to prescribe a singular path forward, but to surface the complex, layered, and interdependent issues that policy actors may wish to explore further. The findings across the consultation suggest that Ireland's higher education system is not starting from scratch, but it does need coordinated, reflective, and inclusive action to ensure that generative AI strengthens, rather than fragments, its educational mission.

7. Conclusion

The integration of generative AI into higher education is not a future challenge, it is a present reality. Across the ten thematic focus groups and the institutional leadership summit conducted for this project, it became clear that staff, students, and leaders are already engaging with AI in diverse, and sometimes conflicting, ways. The questions they are asking are not merely technical, but educational, ethical, and societal: *What do we value in learning? What should students still be able to do themselves? What counts as original work? How can we ensure equity and inclusion in an AI-rich environment?*

This consultation has surfaced both shared concerns and strong appetite for dialogue, clarity, and leadership. While the sector is responding with creativity and commitment, it is doing so in the absence of common frameworks or consistent guidance, at least which had filtered to the participants to date, or protected time and resources for staff to adapt. Participants called not for rigid rules or centralised control, but for strategic coherence, a sense that the sector is moving forward deliberately, ethically, and in alignment with its educational purpose.

At its heart, this report does not present a technical roadmap for AI implementation. Rather, it offers a grounded, system-wide view of how generative AI is currently being encountered by those who teach, support, govern, and learn within Irish higher education. It highlights the tensions and trade-offs that require navigation, and identifies a set of strategic considerations that may support future policymaking and institutional planning.

What is now required is a sustained, collaborative approach, not only to AI tools themselves, but to the kind of higher education Ireland wishes to offer in an era of intelligent technologies. This moment offers an opportunity not only to safeguard what matters in higher education, but to renew and reimagine it. The conversations documented in this report mark an important step toward that goal, and they demonstrate the sector's readiness to engage creatively and conscientiously with the opportunities AI presents. The HEA and higher education institutions can build upon these conversations through continued collaboration, directly informing the forthcoming policy framework to be developed in the coming months.

Appendix 1. Full Thematic Analysis

Enhancement for Teaching: Full Thematic Analysis

Theme 1: AI is Being Used as a Teaching Assistant, But Quietly

Participants described using generative AI to streamline lesson planning, generate examples, rewrite complex text for accessibility, and produce variations of instructional material. These uses were generally small-scale, exploratory, and privately undertaken, without formal institutional recognition or peer discussion.

> *"It saves me time when I'm creating sample questions or rewriting materials for different levels, but I still have to double check everything."*

Some were excited by the idea that AI could reduce repetitive prep work and give staff more time for interaction and feedback. Others were more hesitant, noting that the tools often require correction, and that their outputs can be shallow or misleading.

Pedagogical quality may be compromised if AI-generated materials are used without critical evaluation or adaptation. There is also a risk that teaching becomes less intentional and more reactive if AI begins to shape content without educator input.

Theme 2: Potential to Personalise and Diversify Learning Materials

AI's capacity to generate multiple versions of explanations or assessments was seen as a major opportunity. Participants imagined using it to support differentiated instruction, scaffold learning, and offer alternative formats for students with varying needs.

> *"It gives me a way to respond to the wide range of student needs without doubling my prep time."*

Some participants raised the idea of AI being used to build responsive, interactive learning pathways that adjust to student inputs, though most felt this was aspirational given current capacity.

While the potential for personalisation is real, staff often lack the time, tools, or confidence to build this into their practice. Without institutional support, such innovations remain unevenly distributed and underdeveloped.

Theme 3: Concerns About Deskilling and Erosion of Pedagogical Judgement

Several participants worried that if AI begins to perform core teaching tasks, such as designing content, generating examples, or suggesting pedagogical strategies, staff may lose confidence or skill in these areas over time.

> *"It's seductive, you ask it for help, and it gives you something usable. But the more you use it, the less you reflect on why you teach the way you do."*

This was not just about losing skills, but about diminishing the intellectual and creative role of the educator.

Teaching risks being hollowed out if AI is used as a substitute for pedagogical thinking, rather than a tool to support it.

Theme 4: Uncertainty About Pedagogical Quality of AI Content

AI outputs were described as grammatically fluent but conceptually shallow. Staff expressed concern that AI-generated teaching content often lacked nuance, omitted critical perspectives, or reproduced biases.

> *"It looks great on the surface, but it doesn't really engage with the messiness of knowledge."*

Participants said they lacked clear criteria for judging AI outputs beyond surface-level checks.

Staff require pedagogically grounded frameworks for evaluating AI-generated materials, ideally discipline-specific and co-created with peers.

Theme 5: Lack of Policy, Peer Support, and Institutional Framing

AI use for teaching remains informal, uncoordinated, and largely absent from institutional conversations about curriculum, learning design, or teaching development. This contributes to feelings of isolation.

> *"I'm not sure if what I'm doing is encouraged, ignored, or discouraged, no one's talking about it."*

There was a desire for institutions to acknowledge the growing use of AI, and to offer space for professional discussion and experimentation.

The absence of a shared institutional stance limits innovation and increases risk for staff using AI in their teaching.

Synthesis and Takeaways

- AI is being used to enhance teaching, but its integration is fragmented and under-theorised.
- The challenge is not access to tools, but absence of pedagogical support, reflective space, and shared standards.
- Staff are not asking “Can I use AI?”, but instead “How do I use it well, safely, and meaningfully for students?”

Enhancement for Assessment: Full Thematic Analysis

Theme 1: AI Is Rendering Some Forms of Assessment Obsolete

Staff shared that generative AI tools can now complete many common forms of assessment, including essays, short-answer questions, and reports, to a standard that is difficult to detect or challenge.

> *"What's the point of assigning something that ChatGPT can do in ten seconds?"*

This was seen not just as a challenge to academic integrity, but as a wake-up call to the over-reliance on output-based, text-heavy assessment strategies.

Assessment formats that do not require interpretation, creativity, or personal perspective are increasingly ineffective at distinguishing student learning.

Theme 2: Growing Momentum for Authentic and Process-Oriented Assessment

Rather than respond with detection strategies, participants called for assessments that are harder to outsource, including oral exams, staged submissions, reflective commentaries, and real-world tasks.

> *"It's not about preventing use, it's about making sure the task demands thinking that AI can't replicate."*

There was strong consensus that redesign should foreground student process, interpretation, and ethical decision-making. This would include asking students to explain *how* and *why* they used AI tools, designing assessments with embedded self-reflection and annotation, and using portfolio-based and multi-modal submissions

Theme 3: AI Use in Assessment Requires Transparent and Supportive Policies

Participants expressed discomfort at the lack of consistent messaging around what constitutes acceptable use of AI in coursework. Without clarity, students and staff operate in a climate of uncertainty.

> *"Right now, a student could use AI extensively and still be within the rules, or in breach, depending on who marks it."*

This confusion creates risk for students and contributes to inconsistency in grading and feedback.

Policy should define acceptable practice, with room for disciplinary variation and student co-creation of norms.

Theme 4: Concern About Punitive Responses and Unreliable Detection

There was strong resistance to treating AI use as inherently dishonest. Staff questioned the reliability of detection tools and the ethics of using them without full transparency.

> *"These tools don't explain their reasoning, they just give you a number. That's not good enough for an academic misconduct process."*

Participants worried that punitive responses would harm student trust and disproportionately affect vulnerable learners.

Theme 5: Institutional Structures Inhibit Assessment Innovation

Despite strong support for change, participants described significant structural barriers, including rigid assessment approval processes, lack of time, and fear of inconsistency or student complaint.

> *"We're encouraged to innovate, but punished when things go wrong."*

Several noted that changes to assessment often take a year or more to implement formally, making responsive adaptation difficult.

Synthesis and Takeaways

- Generative AI exposes vulnerabilities in current assessment practice, but also reveals long-standing overdependence on certain formats.
- The solution is not stricter enforcement, but a systemic shift toward assessment that rewards thinking, judgement, and self-awareness.
- Staff are ready to lead this change but need institutional and national support to do so with confidence.

Foundational Skills and Skills Degradation: Full Thematic Analysis

Theme 1: Fear of Superficial Learning and Hollowed-Out Competence

Participants repeatedly expressed concern that students are increasingly using AI to *complete* tasks rather than *learn through* them. As a result, the process of skill acquisition, especially in writing, reasoning, and synthesis, may be bypassed.

> *"They're submitting perfect-looking assignments, but when you talk to them, the thinking just isn't there."*

This led to fears that performance is being confused with competence, and that over time, foundational skills could quietly deteriorate.

Theme 2: Writing as a Cognitive Process, Not Just a Product

Writing was framed as a central academic practice, not only for communication, but for organising thought and developing disciplinary understanding. Participants were concerned that AI tools risk reducing writing to a deliverable, severing it from its pedagogical role.

> *"When students write, they're thinking, structuring, analysing, making connections. If AI does that, what's left?"*

Several participants advocated for preserving writing as an intellectual exercise, even if AI is used in later stages.

Theme 3: Uncritical AI Use Can Undermine Critical Thinking

Staff noted that students often use AI without understanding its limitations, biases, or epistemological stance. This uncritical use may reinforce shallow or generic thinking.

> *"AI can give them a plausible answer, but not a critical one, and they don't always know the difference."*

There was strong support for teaching students *how* to engage critically with AI-generated content, including checking sources, interrogating assumptions, and identifying gaps.

Theme 4: Erosion of Information Literacy and Source Awareness

Participants raised the issue that students are using AI to generate content rather than search for, evaluate, and synthesise evidence from source material. This has implications for information literacy, a foundational academic skill.

> *"Why bother reading papers if you can ask ChatGPT for a summary? But then you miss the nuance, the debate, the evidence."*

Students need guidance on when AI is an appropriate support tool, and when independent research is necessary for academic growth.

Theme 5: The 'De-Skilling' of Educators

Not all concern was student-focused. Some participants reflected on their own growing dependence on AI for generating examples, feedback, or summaries, and questioned what that might mean over time for their own expertise.

> *"I use it to save time, but am I getting worse at the things I used to do well?"*

This mirrored concerns about student skill erosion and pointed to the need for reflective practice among educators themselves.

Synthesis and Takeaways

- Foundational academic skills, especially writing, research, and critical thinking, risk erosion in an AI-rich learning environment unless explicitly protected through curriculum design.
- There is a need to rearticulate what higher education is *for* in the age of generative AI, and which forms of knowledge, skill, and judgement should remain human-developed.
- Institutions can create space for discussion and guidance on how to teach *with* AI while still ensuring students *learn*.

AI Literacy: Full Thematic Analysis

Theme 1: AI Literacy Is Broader Than Tool Proficiency

Participants were clear that AI literacy is not just about knowing how to use tools like ChatGPT or image generators. Instead, it includes understanding their underlying logic, biases, and limitations, and being able to judge when not to use them.

> *"We don't want students who can prompt well, we want students who can think about what the prompt is doing."*

AI literacy includes awareness of algorithmic opacity, bias, hallucination, data provenance, and the political economy of AI.

Institutions should seek to move beyond 'how-to' guides, and support curricula that develop deeper epistemic, ethical, and critical engagement.

Theme 2: Students Are Using AI Without Literacy, and Without Consequence

Participants reported that students are already using AI extensively in their academic work, but often in ways that are unexamined or unacknowledged.

> *"They're not being taught to think about it critically, they're just using it to save time."*

This raises concerns that students are developing dependence on tools without understanding them, or reflecting on how they affect their learning.

Without intervention, AI use will shape learning habits in ways that may be difficult to reverse or regulate.

Theme 3: Staff Are Unprepared to Model or Teach AI Literacy

Many academic staff feel unequipped to engage with AI literacy, especially those in non-technical fields. There was discomfort about 'teaching what I don't fully understand myself.'

> *"I want to teach students how to think critically about AI, but I'm still figuring that out too."*

This lack of staff confidence slows integration and risks reinforcing gaps between disciplines or departments.

Staff development should be pedagogically and ethically framed, not just technical.

Theme 4: Curriculum Integration Is Needed, Not Add-Ons

Participants agreed that AI literacy should be embedded into disciplinary learning, not added as a standalone workshop or optional module.

> *"It has to be relevant to their subject, otherwise it just feels like another digital thing."*

This calls for collaboration between subject experts, learning designers, and digital education specialists.

There is little institutional infrastructure to support this kind of cross-functional curriculum design.

Theme 5: AI Literacy Should Be a Shared, Institution-Wide Responsibility

There was concern that AI literacy is being treated as someone else's job, passed between academic departments, library teams, IT, or careers services.

> *"Everyone sees it as important, but no one sees it as theirs."*

Participants emphasised the need for leadership, coordination, and an institutional narrative that frames AI literacy as a strategic priority.

Synthesis and Takeaways

- AI literacy is a multi-dimensional capability: technical, critical, ethical, and contextual.
- There is a pressing need to define what AI literacy looks like in practice, and who is responsible for delivering it.
- Without institutional coordination and curriculum integration, AI literacy will remain patchy and insufficient.

Ethical AI: Full Thematic Analysis

Theme 1: Ethical Use Is Underdetermined and Highly Contextual

Participants described the current moment as ethically unclear. While many could identify examples of misuse (e.g. full AI-generated assignments), there was no shared understanding of what constituted *acceptable* use in grey areas.

> *"If a student uses it to brainstorm, is that OK? What if they use it to structure the whole essay?"*

What counts as 'cheating' or 'learning support' depends heavily on context, purpose, process, transparency, and student awareness.

No shared norms, vocabulary, or ethical criteria currently exist to support consistent judgment or meaningful conversations.

Theme 2: Lack of Policy Forces Ethical Decision-Making Downward

In the absence of institutional frameworks, participants felt ethical judgments were being pushed to frontline staff, sometimes inconsistently or reluctantly.

> *"I'm making these decisions on the fly, and I have no idea if I'm aligned with my colleagues."*

This creates anxiety, tension across departments, and risk for both students and staff.

Ethical uncertainty should be addressed collectively, not shouldered by individuals operating without support or reference points.

Theme 3: The Tools Themselves Raise Ethical Concerns

Beyond questions of use, several participants raised concerns about the tools themselves: the environmental cost of AI, the biases embedded in training data, the ownership of content, and the exploitation of user input.

> *"We're encouraging students to use tools we don't fully understand, tools built by companies whose interests may not align with education."*

This ethical horizon expands the conversation beyond pedagogy to include social justice, data governance, and corporate influence.

Theme 4: AI Use Can Reinforce or Disrupt Inequities

AI use is not neutral. Participants noted that students with higher digital literacy, confidence, or support were better positioned to benefit from AI. Others may misuse or avoid it due to lack of guidance.

> *"It could help level the playing field, or widen the gap."*

There was a shared sense that unless AI use is framed through an equity lens, existing disadvantages could deepen.

Theme 5: Ethics Requires Dialogue, Not Just Detection

Participants argued strongly against responding to AI use with detection or punishment alone. Instead, they advocated for pedagogical approaches that make ethical reasoning part of the learning process.

> *"We should be helping students think ethically, not just policing them."*

Ethical literacy was seen as both a teaching goal and a process of institutional co-creation.

Synthesis and Takeaways

- Ethical questions about AI use are pervasive, complex, and largely unaddressed by current institutional policies.
- Staff are eager for principled, participatory approaches, not simply compliance models or punitive measures.
- Ethics can be framed as a pedagogical issue as much as a policy one, embedded in how staff and students engage with AI day to day.

Academic Integrity: Full Thematic Analysis

Theme 1: AI Challenges Conventional Definitions of Misconduct

Participants described how traditional conceptions of plagiarism, such as unattributed copying, do not fully capture what is at stake with AI-generated work. For example, if a student uses AI to generate phrasing or structure, is it dishonest, or just poorly guided?

> *"It's not just copying, it's outsourcing thinking. But we don't have a term for that."*

This ambiguity creates inconsistent responses across staff and disciplines, and leaves students confused about expectations.

Definitions of misconduct should evolve to address new forms of assistance and automation.

Theme 2: Over-Reliance on Detection Tools Risks Due Process

Many participants voiced concern about institutions increasingly using AI-detection software, despite known limitations, false positives, and lack of transparency.

> *"We're building disciplinary processes on tools we don't even trust."*

Participants emphasised that reliance on detection undermines the values of fairness and evidence-based adjudication.

Automated detection, used in isolation, risks penalising students without defensible evidence.

Theme 3: Clarity and Transparency Are Missing for Students

Students are using AI tools without understanding what is permitted. Most participants reported that their institutions lacked clear, shared guidelines, leaving responsibility to individual lecturers.

> *"We need to move from don't use AI to here's when and how it's OK."*

This confusion risks accidental misconduct, reinforces inequality, and undermines trust in assessment.

Theme 4: Integrity Is a Culture, Not a Checklist

There was strong consensus that integrity cannot be instilled through rules alone. Instead, it should be embedded in course design, classroom dialogue, and modelling by staff.

> *"If we design out trust, we'll get dishonesty. But if we design for honesty, students will rise to meet it."*

Participants supported moving towards assessment formats that incentivise originality and reflection, rather than punish detection failures.

Theme 5: Staff Feel Exposed and Underequipped

Participants noted that many staff feel uncertain about how to handle suspected misconduct involving AI, especially without legal or institutional clarity.

> *"If I accuse a student, and the tool is wrong, what happens? But if I say nothing, am I complicit?"*

Staff called for guidance, training, and shared decision-making frameworks that go beyond punitive logic.

Synthesis and Takeaways

- Generative AI fundamentally disrupts how academic integrity is defined, demonstrated, and supported.
- Institutions will be more successful in fostering a shared culture of trust, reflection, and clarity rather than just reducing the issue to detection.
- Assessment design, policy communication, and ethical reasoning should be aligned and reinforced across all levels.

Writing: Full Thematic Analysis

Theme 1: Writing Is Being Redefined, But Without Pedagogical Consensus

Participants shared examples of students using AI for brainstorming, drafting, summarising readings, and polishing grammar. While this could support learning, staff worried that AI-generated writing often bypasses deeper intellectual engagement.

> *"Students are handing in fluent work, but it's hollow. They didn't do the heavy lifting."*

This raises questions about whether writing is still a thinking practice, or simply a means of performance.

Theme 2: Authorship and Originality Are Becoming Harder to Define

The concept of authorship is blurring. If a student uses AI to draft an outline, rephrase ideas, or generate transitions, is the final work still their own? Participants found that current frameworks offer little guidance.

> *"We used to talk about plagiarism in terms of sources, now we're talking about where the thinking happened."*

Some staff wanted to allow AI use if transparently disclosed; others were unsure how to judge originality in mixed-authorship texts.

Theme 3: The Writing Process Needs to Be Made Visible Again

There was strong agreement that writing pedagogy should focus more on the process, including planning, revision, and reflection, not just the final product.

> *"We've outsourced too much to the outcome. AI just accelerates that."*

Suggestions included portfolio submissions, staged assignments, and annotated drafts showing student choices and engagement.

Theme 4: Writing Support Services Are Under Pressure and Under-Resourced

Writing centres and learning support staff reported a surge in student questions and confusion about AI use, especially among international students or those unfamiliar with academic writing conventions.

"Students want to use AI to bridge gaps, but they're not getting guidance on how to do that ethically."

This was seen as both a challenge and an opportunity to reorient writing support toward inclusive, dialogic practice.

Theme 5: Staff Need Support to Reimagine Writing in the Curriculum

Many participants felt unprepared to update their writing assignments or provide guidance on acceptable AI use. Some feared being inconsistent, while others worried about undermining legitimate student skill development.

"We've been left to figure it out ourselves. That's not fair, or sustainable."

There was appetite for collaborative curriculum review, co-created guidelines, and examples of good practice.

Synthesis and Takeaways

- Writing in higher education is not just communication, it is a means of learning, thinking, and becoming a disciplinary practitioner.
- Generative AI challenges not only what students write, but how and why they write.
- Reframing writing pedagogy to support process, authorship, and intellectual integrity in an AI-mediated landscape would be beneficial.

Inclusionary Practices: Full Thematic Analysis

Theme 1: Generative AI Can Support Inclusion, If Designed with Purpose

Participants highlighted the potential for AI tools to assist students with dyslexia, non-native English speakers, and those with anxiety or executive functioning challenges. For instance, AI-generated summaries, grammar assistance, or time-saving prompts could enable more equitable participation.

> *"Some students finally feel they can keep up, they have tools that help them express themselves."*

However, participants warned that these benefits depend on proper guidance, transparency, and intentional design, otherwise, they may create dependency or be misused.

Theme 2: Digital Confidence and AI Literacy Are Unevenly Distributed

There was strong concern that students from lower socio-economic backgrounds, mature learners, or those less confident with technology are already falling behind in their use and understanding of AI tools.

> *"There's a growing gap, some students are fluent in AI use, others are terrified of doing something wrong."*

This suggests the need for targeted outreach, just-in-time support, and inclusive training that is co-designed with students.

Theme 3: Institutional AI Strategies Often Omit Inclusion Considerations

Participants described how AI policy and practice are often framed around assessment security or academic innovation, with little or no reference to inclusion or accessibility.

> *"Inclusion is an afterthought, not built into the conversation from the start."*

This omission risks embedding exclusion into AI implementation by default.

Theme 4: AI Can Reinforce Normative Assumptions About Language and Knowledge

Some participants expressed concern that generative AI outputs often reflect dominant cultural norms, including Eurocentric language patterns, assumed background knowledge, and conventional academic structures.

> *"The writing it produces is grammatically perfect, but it's not their voice. That matters."*

This can subtly erase the linguistic, cultural, or epistemic diversity students bring, especially for multilingual or international learners.

Theme 5: Equity Requires Structural, Not Just Technical, Solutions

While AI tools may assist individual learners, participants stressed that inclusion requires systemic commitment, including curriculum review, staff development, and resource allocation.

> *"We can't just hand out tools and say we've done inclusion. It has to be a whole-institution commitment."*

Participants called for the integration of AI considerations into Universal Design for Learning, access planning, and widening participation strategies.

Synthesis and Takeaways

- Generative AI presents both risk and opportunity for inclusion, neither outcome is guaranteed.
- Institutions should make inclusion a first-order concern in AI planning, not a retrofit or exception.
- Inclusion-focused strategies go beyond access to tools, instead addressing student voice, curriculum design, and the legitimacy of diverse ways of knowing and expressing.

Infrastructure: Full Thematic Analysis

Theme 1: Infrastructure Is Inadequate and Unevenly Developed

Participants described a patchwork of infrastructure across institutions, with some having exploratory licences, others relying on free versions, and many without any guidance at all.

> *"People are using whatever's available, sometimes tools that haven't been vetted or approved."*

This environment creates inconsistency, increases data privacy risks, and reinforces digital inequity.

Theme 2: Governance and Procurement Are Unclear or Missing

Participants repeatedly noted the absence of clear governance processes around which AI tools could be adopted, who was responsible for oversight, and what ethical standards were applied.

> *"We've no idea who is evaluating these tools, or what criteria they're using."*

Some were concerned about commercial tools entering classrooms through informal channels, bypassing procurement or data protection protocols.

Theme 3: Tool Adoption Is Often Driven by Individuals, Not Strategy

Much of the AI activity described was led by enthusiastic individuals or teams, rather than institutionally endorsed platforms or initiatives.

> *"It's all bottom-up, which is good for innovation, but risky without support or coordination."*

This raises questions about sustainability, quality assurance, and alignment with teaching goals.

Theme 4: Infrastructure Includes Human and Organisational Capacity

Participants emphasised that infrastructure is not just about licences or platforms, it also includes support staff, professional development, and time.

> *"You can't roll out a new tool and expect staff to know how to use it pedagogically."*

There was strong support for investing in cross-functional collaboration between IT, academic developers, and teaching staff.

Theme 5: Infrastructure Benefits from Embedded Values and Standards

The group warned that platform choices embed assumptions about pedagogy, privacy, and power. Decisions about which tools are licensed should be made with reference to educational values, not just functionality or cost.

> *"Procurement is pedagogical. It's not neutral."*

Concerns included terms of service, surveillance, commercial use of student data, and alignment with academic freedom.

Synthesis and Takeaways

- Generative AI infrastructure is underdeveloped across Irish higher education, posing risks to equity, security, and educational coherence.
- Infrastructure is best understood as a whole system, including tools, policies, people, and processes.
- National coordination and investment are needed to support scalable, ethical, and pedagogically aligned AI infrastructure.

SETL Committee: Full Thematic Analysis

Theme 1: Reaffirming Pedagogical Purpose in an AI Context

Participants emphasised that AI use in education is best guided by a clear sense of pedagogical purpose. The proliferation of AI tools risks pushing institutions into reactive positions, where teaching becomes tool-led rather than value-led.

> *"If we don't decide what kind of learning matters, AI will decide for us."*

Participants argued for renewed attention to educational values, disciplinary identity, and the social function of higher education.

Theme 2: Student Engagement Risks Erosion Without Thoughtful Design

There was concern that if AI is used to automate or shortcut learning tasks, students may become passive recipients rather than active participants.

> *"We need to give students reasons to engage, not just punish them for disengaging."*

Engagement strategies can evolve to include meaningful challenges, social learning, and co-authorship of learning goals.

Theme 3: Flexible, Authentic, and Future-Facing Curriculum Models

Participants agreed that AI should not be treated as an external module or problem to solve, but as a catalyst for curriculum renewal. This includes assessment redesign, authentic learning, and interdisciplinarity.

> *"AI is not a topic, it's a condition. It affects how we think, write, assess, and relate."*

Traditional programme structures were seen as too rigid and slow-moving to adapt meaningfully to the pace of change.

Theme 4: Staff Capacity Is the Limiting Factor in Responsible AI Adoption

Many staff are expected to make rapid pedagogical shifts with little training, time, or support. Participants stressed that digital and pedagogical capacity-building go hand-in-hand.

> *"We're asking staff to do something transformative, but without the tools or headspace to do it well."*

There was strong support for institution-wide professional learning strategies rooted in collaboration, reflection, and disciplinary relevance.

Theme 5: Students As Partners in Shaping AI-Enhanced Education

Rather than treat students as compliance risks or assessment problems, participants advocated for a partnership model, involving students in co-design, feedback, and innovation.

> *"If we build this without them, we're not preparing them for the world they're already in."*

Partnership was also seen as a route to more ethical, equitable, and trusted educational innovation.

Synthesis and Takeaways

- Generative AI should be integrated into long-term strategies for educational transformation, not treated as a compliance or technology issue alone.
- Student engagement, teaching design, and curriculum models will evolve, with AI as one of many drivers of change.
- Leadership, staff development, and participatory culture are key enablers for success.

Institutional Leadership Perspectives: Full Thematic Analysis

Theme 1: AI Requires a Rethinking of Institutional Purpose and Values

Leaders voiced the need to re-express what higher education is for in an AI-enabled world. With automation accelerating, institutions will be faced with decisions around the protection and redefinition of the human, ethical, and intellectual dimensions of learning.

> *"By 2030, every graduate in Ireland could be entering a workforce transformed by AI."*

Key values cited for defence and adaptation included: integrity, human dignity, critical thinking, community, and service. There was agreement that 'integrity' should be redefined to reflect the nature of process-oriented, AI-assisted learning.

Theme 2: Governance is Fragmented, Reactive, and Losing Credibility

Participants described a landscape of conflicting local policies, inconsistent practices, and unclear leadership mandates.

> *"Staff are being told to teach ethical AI, but the tools themselves may go against their values."*

This lack of clarity is eroding trust. Students perceive double standards, while staff express fatigue and confusion. Leaders called for specific, transparent, and actionable frameworks, not vague or overly broad guidelines.

Theme 3: The Sector Is Exposed to Strategic and Reputational Risk

Institutions face multiple vectors of risk, from unethical grading and copyright breaches to disengagement and diminished graduate competence.

> *"Students are getting degrees without engaging. That is a reputational time bomb."*

Leaders noted that public, employer, and governmental trust may be affected if institutions are seen to neglect rigour, accountability, or quality in the face of AI.

Theme 4: Institutional Readiness is About Culture, Not Just Tools

Participants strongly rejected the notion that being "AI-ready" is purely a matter of technology adoption. They defined readiness as a combination of:

- Clarity of institutional values
- Academic integrity frameworks
- Student engagement strategies
- Staff training and digital literacy
- Open and consistent communication
- Capacity for ethical leadership and agile governance

"Academic integrity and student engagement matter most. Leadership alignment? Not so much, because trust in leadership is low."

This gap between institutional aspiration and perceived leadership capability was a recurring concern, indicating a need for more transparent communication and inclusive planning.

Theme 5: Adaptive, National Coordination Is Urgently Needed

Leaders advocated for a shared national framework that allows for local adaptation, not rigid policies but "practice principles".

> *"We shouldn't call them policies. AI changes too fast. We need something more dynamic."*

Examples from Scotland and Hong Kong were cited where universities had pooled resources, standards, and toolkits. There was a strong desire for shared infrastructure (e.g. AI literacy content, copyright guidance, accessibility standards), delivered through an equitable, sector-led approach.

Theme 6: Institutions Can Choose to Lead, Not Just Respond

Participants warned against treating AI as an external disruption to be survived. Instead, they urged institutions to shape the discourse, embedding critical thinking, ethical literacy, and social responsibility into AI strategies.

> *"The time has passed to tell students not to use AI. We are challenged with teaching them how to use it ethically, and why."*

This included calls for interdisciplinary curriculum reform, lifelong staff development, and sustained student-staff dialogue.

Synthesis and Takeaways

- Generative AI is not only a pedagogical issue, it is a governance, reputational, and cultural challenge.
- Institutional leadership should define clear principles, steward difficult conversations, and model ethical action.
- The sector has an opportunity to lead internationally by adopting a collaborative, principled, and adaptive stance.

Appendix 2. List of Participants

The consultation process engaged more than 80 individuals across the Irish higher education system. Two complementary strands were undertaken: focus groups with staff and students, and a leadership consultation with senior institutional representatives. Together, these activities ensured that the perspectives of a broad cross-section of the sector were included.

Institutional Representation

Participants were primarily drawn from across each of the HEA-funded higher education institutions:

Atlantic Technological University
Dublin City University
Dun Laoghaire Institute of Art and Design & Technology
Dundalk Institute of Technology
Mary Immaculate College
Maynooth University
Munster Technological University
National College of Art & Design
RCSI University of Medicine and Health Sciences
Royal Irish Academy
South East Technological University
St Angela's College
Technological University Dublin
Technological University Shannon: Midlands Midwest
Trinity College Dublin
University College Cork
University College Dublin
University of Galway
University of Limerick

Stakeholder Representation

In addition to HEIs, the process included representatives from key stakeholders, including:

Department of Further and Higher Education, Research, Innovation and Science
Union of Students in Ireland
Government of Ireland AI Advisory Council
National Academic Integrity Network
HEAnet
AHEAD
Quality and Qualifications Ireland
Enterprise Ireland
Microsoft