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

Generative AI in Higher Education Teaching & Learning

AI Literacy Training

Contributors

James O'Sullivan
Colin Lowry
Ross Woods
Tim Conlon

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

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

Version 1.0, December 2025
Higher Education Authority, Dublin
DOI: 10.82110/3sxr-z770

How to cite:
O'Sullivan, James, Colin Lowry, Ross Woods & Tim Conlon. *Generative AI in Higher Education Teaching & Learning: AI Literacy Training*. Dublin: Higher Education Authority, 2025. DOI: 10.82110/3sxr-z770.

These recommendations propose baseline requirements for AI literacy programmes within higher education institutions. The objective is to ensure that all members of the academic community possess sufficient understanding of generative AI systems to use them responsibly, recognise their limitations, and comply with institutional policy.

AI literacy training must address both technical and ethical dimensions. This dual emphasis is essential, as without technical understanding, ethical guidance becomes abstract. Staff cannot make sound judgements about when to trust an AI-generated output if they do not understand why such outputs are sometimes fluent yet wrong, while equally, technical knowledge without ethical grounding risks producing users who are skilled at extracting outputs but indifferent to the consequences of their use.

The technical component of AI literacy serves a further strategic purpose. When staff understand how large language models work, including how they are trained, why they hallucinate, and what techniques mitigate their weaknesses, they are equipped to develop discipline-specific guidance without relying on central resources for every decision. Technical literacy is the foundation upon which contextualised, discipline-appropriate practice can be built.

This document does not prescribe a single curriculum because institutions differ in scale, mission, and disciplinary composition. What it does prescribe is a set of outcomes, principles, and structural requirements that any effective programme should aim to meet. The aim is a standard wherein participants can identify the main types of AI systems, explain how generative systems produce outputs in plain terms, and apply ethical principles such as proportionality, transparency, accessibility, and continuous improvement to their daily work.

Training must be rooted in institutional policy so that staff, as learners, engage with AI in more than the abstract sense but with the precise responsibilities and risks that govern their context.

Developing AI literacy across institutions has implications for capacity, capability, and ongoing professional development. Generative AI tools are already embedded in the workflows of students and staff, whether institutions sanction their use or not. Untrained users make predictable errors, trusting outputs that should be verified, disclosing data to systems that harvest it for training, and failing to recognise when AI-generated content violates accessibility requirements or disciplinary norms. These errors carry costs to academic integrity and to the quality of student learning, and in turn, to institutional

reputation.

AI literacy development supports both risk mitigation and institutional capacity building. Institutions that equip their staff with genuine technical understanding create a distributed resource: academics who can advise colleagues, develop discipline-specific exemplars, and respond to new tools as they emerge without waiting for central guidance. This distributed model is the only sustainable approach in a domain where tools and practices evolve faster than any central team can track.

Approaches that focus narrowly on compliance, without supporting deeper understanding, may limit staff capacity to adapt as tools and contexts evolve. Where training is perceived as disconnected from disciplinary practice, engagement may be reduced and AI literacy treated as a procedural requirement rather than a professional development opportunity.

Effective AI literacy development has implications for staffing, infrastructure, and ongoing maintenance. This may include access to subject-matter expertise with both technical and pedagogical understanding; platforms capable of hosting and tracking training activity; regular review cycles to ensure accuracy and relevance; opportunities for discipline-specific contextualisation; and structured involvement of students in co-design and peer support roles. Institutions will need to consider how such elements can be supported within existing capacity and delivery models.

The scale and form of support for AI literacy development will vary by institution, reflecting differing contexts and capacities. However, effective provision depends on clear alignment between expectations, responsibility, and available capacity. Where responsibility for AI literacy is assigned without adequate support, there is a risk that provision may be inconsistent or insufficient to achieve intended outcomes.

Governance and accountability

Governance of AI literacy requires clear lines of accountability combined with active community partnership. Each institution should identify an Academic Owner responsible for ensuring the intellectual integrity and relevance of training content. This role requires subject-matter expertise in generative AI systems and their application in higher education, and should be held by staff with appropriate expertise in academic practice and pedagogy. Alongside the Academic Owner, a Service Owner should be responsible for delivery mechanics, record-keeping, accessibility compliance, and platform maintenance. Senior

institutional oversight, typically at registrar or equivalent level, supports assurance that training is mandatory where required, adequately resourced, and subject to regular review.

Partnership with students strengthens both the relevance and the legitimacy of training. Institutions should establish Student AI Ambassador programmes in which students assist in the co-design of student-facing scenarios and the peer-delivery of orientation materials. Students use AI tools in ways that staff may not anticipate so their involvement ensures that training reflects actual usage patterns rather than staff assumptions about how tools are employed. A designated public contact point should be maintained to allow staff and students to seek clarification on roles, resources, and policy interpretation.

Audience and training

Training delivery should be stratified to ensure every member of the community possesses knowledge appropriate to their responsibilities. A single, undifferentiated programme cannot serve both a first-year undergraduate encountering AI for the first time and a programme director responsible for assessment design across multiple modules.

All students should complete a concise orientation, normally delivered within the first weeks of study and prior to engaging in any assessed work where generative AI use may be relevant. This orientation should introduce core concepts without assuming prior technical knowledge, and completion should be recorded and verifiable by programme teams.

All staff should complete a core literacy course. This course should be more substantial than the student orientation, reflecting staff members' broader responsibilities and their role in advising students. Staff holding high-impact responsibilities require additional, role-specific training. Those involved in assessment design must understand how AI tools interact with different assessment formats and how to write briefs that clarify permitted and prohibited use. Those involved in tool approval or procurement should understand data protection implications, vendor incentives, and the distinction between institutional instances and public-facing tools. Those handling student data in conjunction with AI tools must understand when data protection impact assessments are required and how to request them.

Senior leaders and approvers require a distinct briefing focused on oversight governance, risk appetite, escalation protocols, and their responsibilities under institutional and

regulatory frameworks. This briefing need not cover technical detail at the same depth as the staff course, but it must ensure that decision-makers understand the stakes of the decisions they are asked to make.

Upon successful completion, participants should be able to classify the principal types of AI systems in use at the institution and describe the risks and safeguards associated with each. They should be able to explain, in ordinary language, how large language models are trained, why they produce outputs that are fluent but sometimes inaccurate, including so-called hallucinations, and how techniques such as structured prompting and retrieval-augmented generation can mitigate these risks.

Beyond technical mechanics, participants should demonstrate critical vendor literacy, meaning they can identify commercial incentives behind AI tools, evaluate terms of service to determine whether user data is protected or harvested for model training, and distinguish between institutional deployments with negotiated data protections and public-facing tools without such protections.

Participants should be able to apply institutional principles in practice, deciding what constitutes proportionate oversight for a given task, executing transparent disclosures when AI has been used, ensuring that outputs meet accessibility standards or that equivalent alternatives are provided, and maintaining logs that facilitate continuous improvement and incident review.

Course content

Effective training should address three substantive, interdependent areas: the types of AI, the mechanics of generative AI, and ethical application. Understanding types of AI provides the vocabulary to distinguish generative systems from other tools, while understanding mechanics explains why generative systems behave as they do. Understanding ethics provides the framework for deciding when and how to use them, and enables staff to move from uncritical to critical forms of adoption and rejection.

Types of AI

Participants should be introduced to the main families of AI, including rule-based systems that follow explicit programmed logic, classical machine learning systems that identify patterns in data, generative systems that produce new content, and compound workflows that combine multiple approaches. These concepts should be grounded in higher education use cases, which will help participants recognise AI where they encounter it and understand why different systems carry different risks.

Generative AI and large language models

Training should focus on developing a working understanding of how large language models generate outputs, so that staff can make informed, discipline-specific judgements about when and how such systems may be appropriate to use. Rather than treating AI as a black box to be accepted or rejected in the abstract, training should make underlying mechanisms visible through concrete examples of documented failure. A curated repository of real cases, including hallucinated facts, fabricated citations, and confidently incorrect analyses, allows participants to see how and why these systems break down in particular contexts. Such examples, whether drawn from invented legal precedents, plausible-sounding but non-existent scholarly references, or misleading summaries of source material, enable staff to determine for themselves where verification is indispensable, where use may be defensible, and where deployment is simply incompatible with disciplinary standards.

Ethical use and data sovereignty

The ethical component of training is anchored in the institution's stated principles. Participants should examine what proportionality means in practice, distinguishing between low-stakes formative support where AI assistance may be freely permitted and high-stakes summative assessment where it may be restricted or prohibited.

Continuous improvement should be framed as a habit of practice which might entail logging AI interactions and reviewing outcomes, as well as converting mistakes into safer workflows. This is not bureaucratic record-keeping for its own sake but the mechanism by which individuals and institutions learn from experience.

Data sovereignty requires particular attention as participants should understand their rights regarding input data, including what happens to the text they enter into an AI tool, whether it may be used to train future models, and how this varies between tools. They should be able to distinguish between institutional instances of tools, where data handling terms have been negotiated, and public versions that may harvest inputs for training. They should understand when entering data into an AI tool may constitute a data protection breach and how to avoid such breaches.

A further dimension of this training concerns the wider geopolitical context in which generative AI systems are owned and governed. Participants should be encouraged to situate everyday pedagogical and research uses of AI within a landscape shaped by transnational corporations, uneven regulatory regimes, and intensifying competition over data, compute, and linguistic resources. Understanding where systems are built, whose legal jurisdictions they fall under, and which languages and cultures they privilege is a practical concern, bearing directly on questions of academic independence, public trust, and institutional risk. In an evolving geopolitical and regulatory context, informed restraint and critical awareness become part of ethical practice, ensuring that institutional reliance on AI tools does not inadvertently entrench external dependencies or compromise the values of a publicly accountable higher education system.

Discipline-specific AI literacy training

Generic awareness is insufficient because what constitutes appropriate, proportionate, or transparent use of generative AI is inseparable from disciplinary method, assessment form, and the nature of the materials being worked with. While core training can establish a shared technical and ethical baseline, it is only through disciplinary translation that this foundation becomes meaningful. Faculties and schools should interpret institutional principles in light of their own epistemic standards and pedagogical practices.

Each discipline may develop concrete exemplars that demonstrate how institutional expectations apply in practice. In the humanities, this may involve clarifying where generative tools can support preparatory activities such as idea generation or structural planning, and where their use would undermine the interpretive labour that defines the field. In the sciences, exemplars might show legitimate applications in data handling or code optimisation, while making explicit that the fabrication, manipulation, or selective

presentation of results remains misconduct irrespective of the technology involved. In professionally accredited programmes such as medicine, law, or teacher education, exemplars must address regulatory obligations, professional ethics, and the heightened risks associated with AI use in contexts where error carries real-world consequences.

Training should therefore engage with worked examples that are recognisable to each discipline: drafting assessment briefs that articulate permissible AI use with clarity, evaluating AI-generated material against disciplinary standards of evidence, and identifying ethical, accessibility, or bias-related concerns specific to the field. These exemplars should be shared across the institution, allowing programmes to adapt them rather than reinventing guidance in isolation.

The aim is to ensure that staff and students are able to exercise informed judgement by applying institutional principles to the realities of their own field. Requiring each school or programme to contribute a small number of exemplars within an initial adoption period, and to revisit them as tools and practices evolve, embeds this judgement as a living part of academic practice rather than a one-off compliance exercise.

Delivery, assessment, and integration

Training should be short enough to complete without undue burden and accessible to all participants regardless of disability or technological constraint. Delivery may combine self-paced online resources with optional live clinics for those who benefit from interactive explanation. Training should be available on demand for new staff and students at any point in the year, not restricted to fixed intake periods.

Materials should be provided in accessible formats by default, and where a tool or platform used in training is not accessible to a participant, an equivalent route must be provided without disadvantage. This is a legal requirement, but it is also a matter of institutional credibility: training that emphasises accessibility while being delivered inaccessibly undermines its own message.

Content should be refreshed in response to major advances in technology or when significant incidents occur. The technology landscape shifts quickly and training based on tools or interfaces that no longer exist wastes participants' time and erodes trust in the

programme.

Participants should complete assessment to demonstrate understanding of the three substantive areas, with completion records maintained in the learning platform and made available to programme teams and quality assurance processes. Records should be retained according to institutional schedules and used primarily for support and improvement. Where completion gaps are identified, the appropriate response is additional support, not sanction, at least until reasonable opportunity to complete has been provided.

AI literacy training should be integrated with institutional processes and not treated as a standalone module disconnected from practice. Participants should be directed to the specific policy sections that govern assessment, to the mechanism for disclosing and recording AI use, to the process for tool approval and exceptions, and to incident and harm response procedures. They should be shown how to seek advice on data protection, including when a data protection impact assessment is required, and how to request accessibility support.

This integration serves two purposes, ensuring that training is practically useful rather than just informative, and it positions AI literacy as part of the institution's broader framework of compliance and governance rather than an isolated initiative that participants may forget once completed.

Quality assurance, review, and communication

The Academic Owner should check training content regularly for factual accuracy and for completeness against evolving institutional and regulatory requirements. Updates should incorporate lessons from incidents within the institution, feedback from participants, and developments in the broader field.

The Service Owner should maintain a change-log so that staff and students can see what has been updated and why. Where a significant change affects practice, targeted refresher communications are required rather than expecting participants to discover changes by re-taking the full training.

Institutions should support training with plain-language communication. A short summary for students and staff, published alongside the training, should explain what is expected,

what support is available, and how to ask questions. The tone of these communications matters, as the aim is to build confidence and competence.

If institutions expect AI to be used responsibly, it must also commits to providing the knowledge and support that make responsible use achievable.