



Co-funded by
the European Union

ADMIT

Recommendations for EU policy on the ethical use of LLM/generative AI in higher education



The research reported in this document was carried out in the ADMITE EU funded project. Proposal ID 101134520 under the call ERASMUS-2023-PCOOP-ENGO

Conti, L., & Breda, M. (2026). Recommendations for EU policy on the ethical use of LLM/generative AI in higher education. Zenodo. [DOI: 10.5281/zenodo.23237363](https://doi.org/10.5281/zenodo.23237363)

License used: This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License: <https://creativecommons.org/licenses/by-sa/4.0/>.

ShareAlike: If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.



Co-funded by
the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.

Executive Summary

The rapid diffusion of Large Language Models (LLMs) and Generative AI is reshaping how higher education institutions organise teaching, learning, assessment, student support and academic governance. This deliverable addresses the use of LLMs and Generative AI in the educational activities of higher education - teaching, learning, assessment, student support and academic governance - rather than in scientific research. Increasingly embedded in learning platforms, tutoring services, feedback processes, administrative support, assessment workflows and decision-making environments, these technologies offer opportunities for innovation, personalisation, accessibility and efficiency, while raising risks for academic integrity, learning, student rights, transparency, accountability, equity, data governance and quality assurance. These dynamics take on specific characteristics in distance higher education institutions such as the universities that deliver all or most of their teaching through online platforms and whose student body is geographically dispersed. Such institutions, which exist in varying legal and regulatory forms across EU Member States, present a distinct combination of opportunities and constraints: GenAI is already deeply embedded in their digital learning environment, and their assessment and quality assurance arrangements may differ substantially from those of campus-based universities. EU policy guidance should explicitly recognise this diversity of institutional delivery models.

European policy already offers important instruments for addressing these challenges. The EU AI Act provides the legal framework for risk-based classification, high-risk AI systems and general-purpose AI models; the Digital Education Action Plan supports digital education transformation; and the Standards and Guidelines for Quality Assurance in the European Higher Education Area (ESG 2015) provide the recognised framework for higher education quality assurance. However, these instruments do not yet form a common European operational framework specifically for LLMs and Generative AI in higher education. The ADMIT outputs address this gap by translating existing principles into the academic context. The ethical taxonomy, the guide for the ethical use of LLM / Generative AI systems, and the institutional self-assessment tool offer higher education-specific guidance on transparency, accountability, academic integrity, the safeguarding of learning, human oversight, data governance, fairness, accessibility and quality assurance. They do not replace European regulation or policy, but clarify how it can be interpreted and applied in universities. The central issue is therefore not a lack of regulation, but the need to translate existing European frameworks into practical institutional guidance on classification, disclosure, academic integrity, contestability, procurement, data governance, equitable access, opt-out safeguards and quality assurance.

A European policy layer for higher education would help manage this complexity without imposing a separate regulatory regime. The recommendations proposed in this section are designed to support trustworthy innovation. They do not aim to restrict the use of Generative AI in higher education. Rather, they aim to ensure that its adoption remains pedagogically meaningful, legally compliant, ethically responsible, socially inclusive and institutionally accountable. This requires guidance, templates, disclosure standards, classification tools, procurement clauses, academic integrity standards, AI literacy resources, self-assessment instruments and monitoring mechanisms.

The section also highlights the importance of avoiding fragmentation. As national implementation initiatives begin to translate the AI Act into domestic legal and institutional arrangements, higher education institutions across Europe may face divergent expectations, inconsistent safeguards and uneven levels of preparedness. A common European reference framework for higher education would support coherence while respecting institutional autonomy and national legal contexts. Overall, the section argues for a practical and proportionate European approach: one that connects AI regulation with academic values, student protection, teacher responsibility, institutional governance and quality assurance. The objective is to ensure that LLMs and Generative AI systems are not simply adopted as technological tools, but governed as part of the academic mission of higher education.

Key Recommendations at a Glance

For Students: Clear disclosure when AI is used in teaching, assessment or support; transparent rules on permitted and prohibited GenAI uses; accessible appeal and opt-out mechanisms for AI-mediated decisions; equitable access to GenAI tools in formal learning activities; structured AI literacy as part of academic practice; learning activities and curricula adapted to the responsible integration of GenAI; protection from unfair or automated misconduct accusations.

For Teachers: Operational guidance on approved tools and permitted data inputs; training on AI output verification, hallucination risks, bias and academic integrity (as part of a shared critical AI literacy common to students and teachers); support for assessment redesign in GenAI-rich environments; meaningful human oversight responsibilities in AI-assisted grading, feedback and proctoring; professional development connecting AI literacy with pedagogical practice.

For Institutions (including administrative and support layers): Risk-based classification of all GenAI deployments; documented governance framework with clear responsibility assignment; procurement standards with data governance, cybersecurity and accessibility requirements; periodic self-assessment using the ADMIT taxonomy and Institutional Self-Assessment Tool; integration of GenAI governance into ESG-aligned quality assurance; institutional academic integrity policies with disclosure norms, assessment redesign and fair misconduct procedures; EU-level monitoring and iterative policy updating.

Table of Contents

Executive Summary	4
Table of Contents.....	6
1 Introduction.....	8
2 Classification of Educational LLMs under the AI Act	9
2.1 From general-purpose model to educational AI system	9
2.2 Risk tiers and educational use cases	10
2.3 Applicability of Annex III to higher education	12
2.4 General-purpose AI obligations and the higher education deployment context	13
2.5 Cross-reference with the ADMIT ethical taxonomy	15
2.6 Implications for policy recommendations	18
3 Current Regulatory, Policy and Implementation Gaps in Relation to Higher Education	19
3.1 The AI Act does not define a sector-specific governance model for universities.....	19
3.2 Academic integrity is not addressed as a distinct regulatory category	20
3.3 Existing EU frameworks do not provide sufficient transparency requirements for educational GenAI use.....	21
3.4 The provider-deployer distinction leaves unresolved accountability questions	23
3.5 Human oversight and appeal mechanisms are not specified for academic contexts	24
3.6 The Digital Education Action Plan does not yet address the specific governance challenges of LLMs	26
3.7 Equity, accessibility and opt-out mechanisms remain insufficiently regulated	27
3.8 Data governance and procurement standards are not sufficiently adapted to educational LLM deployment.....	29
3.9 GenAI governance is not yet embedded in European quality assurance mechanisms	31
3.10 Lack of systematic monitoring and iterative policy updating.....	34
3.11 Synthesis: the need for a sector-specific EU policy layer	35
4 Concrete Recommendations to the European Commission.....	40
4.1 Establish EU-level guidance for the classification and deployment of LLM-based systems in higher education	41
4.2 Define a common transparency standard for educational GenAI use.....	42
4.3 Establish minimum safeguards for AI-mediated academic decisions.....	43
4.4 Develop common European reference principles for institutional academic integrity policies on GenAI.....	44

4.5	Ensure equitable access, accessibility and opt-out safeguards	45
4.6	Ensure environmentally proportionate GenAI deployment	46
4.7	Make structured AI literacy and teacher training a precondition for formal institutional deployment	46
4.8	Introduce periodic institutional self-assessment and EU-level monitoring.....	47
4.9	Embed GenAI governance into ESG-aligned quality assurance.....	48
4.10	Support a European implementation route based on guidance, templates and shared tools 49	
	References	54
	Annex A - Policy Implementation Toolkit: Operational Templates.....	56
	A.1 Classification Checklist Template	56
	A.2 Disclosure Notice Template	56
	A.3 Procurement Clause Model.....	57
	A.4 Academic Integrity Policy Template.....	57
	A.5 Self-Assessment Cycle Guide.....	58
	A.6 Appeal Procedure Template	58

1 Introduction

This section formulates policy recommendations for the ethical and responsible use of LLMs and Generative AI systems in higher education. The analysis takes into account the results produced within ADMIT WP5 as a significant point of reference, with particular attention to the ethical taxonomy developed in Task T5.2 (Carbonel et al., 2024), the Institutional Self-Assessment Tool developed in Task T5.4 (Triantafyllopoulos et al., 2025a), and the Guide of ethical use of LLM / Generative AI systems in Higher Education developed in Task T5.3 (Triantafyllopoulos et al., 2025b). The purpose of this section is to connect the ethical, pedagogical and institutional issues identified by ADMIT with the current European regulatory and policy landscape. The analysis focuses primarily on the EU AI Act, Regulation (EU) 2024/1689 (European Union, 2024), with specific attention to the risk-based classification of AI systems, the provisions on general-purpose AI models, and the high-risk categories applicable to education and vocational training. It also considers the Digital Education Action Plan 2021–2027 (European Commission, 2020) as the broader EU policy framework for digital education, and the ESG 2015 (ENQA et al., 2015) as the reference framework for quality assurance in the European Higher Education Area.

For the purposes of this deliverable, the following terminology applies. Generative AI refers to AI systems capable of generating new content, whereas LLMs constitute one category of models that may underpin such systems. A GPAI model denotes the model-level category established by the AI Act, while an LLM-based or GenAI-enabled system refers to a specific downstream application deployed for an educational or administrative purpose. The applicable governance requirements depend primarily on the intended purpose and institutional use of the system rather than on the underlying model alone.

The European governance landscape for AI and digital education is already structured around several complementary instruments, each addressing a different layer of the problem: legal compliance, digital transformation, ethical responsibility and quality assurance. At the same time, the AI Act is now moving from the level of European regulation to concrete implementation, including through national measures that translate its principles into domestic institutional arrangements. For example, the Italian Council of Ministers Press Release No. 177 of 10 June 2026 presents the national implementation package on artificial intelligence as a set of measures intended to ensure the application of Regulation (EU) 2024/1689 in the Italian legal order, rather than as an alternative regulatory regime (Presidency of the Council of Ministers, 2026). This example illustrates a broader dynamic: national implementation initiatives are beginning to translate the AI Act into domestic legal and institutional frameworks, introducing provisions on compulsory AI literacy in higher education and on the requirement that consequential decisions affecting individuals must not be taken exclusively by automated systems. These nationally enacted principles reflect the spirit of the AI Act and are relevant to any EU-level guidance for higher education. However, the primacy of the European regulatory framework must be preserved: national measures may offer useful implementation models, but they cannot set the standard for guidance addressed to the European Higher Education Area as a whole. Against this background, higher education requires a dedicated European coordination layer, able to translate the AI Act and related policy instruments into shared

criteria for academic governance, and to reduce the risk of uneven interpretations across national implementation processes and institutional practices. The AI Act establishes the legal baseline, although it does not define in detail how universities should govern LLMs in teaching, learning, assessment, student support, academic integrity procedures, procurement, institutional accountability and quality assurance. Similarly, the Digital Education Action Plan promotes high-quality, inclusive and accessible digital education, but it does not provide a detailed operational framework for the current generation of LLM-based systems. The ESG 2015 offer a strong institutional channel for quality assurance, which should now be interpreted in light of the new challenges introduced by Generative AI.

On this basis, after this brief introduction (§1), the deliverable proceeds in three steps. First (§2), it analyses how educational uses of LLMs should be classified under the EU AI Act, distinguishing between general-purpose models, low-impact assistive uses, transparency-sensitive applications and high-risk educational use cases. Second (§3), it identifies the main regulatory and policy gaps that remain unresolved in relation to higher education. Third (§4), it formulates concrete recommendations addressed to the European Commission, aimed at supporting a coherent, transparent, accountable and human-centred European approach to the use of LLMs and Generative AI systems in higher education.

2 Classification of Educational LLMs under the AI Act

The deployment of LLMs and Generative AI systems in higher education must be assessed within the EU AI Act's risk-based framework (European Union, 2024). Classification cannot depend on the technology alone: in universities, LLMs may support activities ranging from optional study assistance to automated assessment, student monitoring, academic integrity procedures, admission support, tutoring, administrative decision-making or institutional analytics. These uses raise different legal and ethical concerns.

The key question is therefore not whether LLMs are high-risk as such, but under which institutional conditions and educational purposes they may fall within AI Act categories. A model used for brainstorming, translation, language improvement or self-directed study differs from one used to evaluate learning outcomes, support grading, detect prohibited behaviour, rank applicants, recommend admission, monitor progression or influence access to educational opportunities. Risk must be interpreted contextually, considering intended purpose, deployment setting, autonomy, individual impact and effects on fundamental rights.

2.1 From general-purpose model to educational AI system

LLMs are commonly developed as general-purpose AI models (GPAI), meaning they are not designed for a single educational function. The same model can generate text, code, summaries, explanations, feedback, translations or recommendations, and can be integrated into many downstream applications. In higher education, one model may therefore serve as a student-support

chatbot, writing assistant, teaching-material generator, formative feedback tool, assessment support system or component of a broader decision-making workflow.

It is therefore essential to distinguish the model from the educational system or service in which it is used. The model is the underlying general-purpose technology; the educational AI system is the concrete application deployed in a specific academic context, for a defined purpose and with consequences for students, teachers or applicants. A GPAI model becomes relevant for higher education governance when embedded in institutional processes such as teaching, learning support, assessment, admission, proctoring, academic integrity, student support or administration. The AI Act reflects this distinction by differentiating between the provider of a GPAI model and the deployer of an AI system (European Union, 2024). Universities will usually procure, license, configure, adapt or integrate third-party models rather than develop them. They are not responsible for the whole model-development process, but they are responsible for the system's institutional use. They must therefore assess not only what the model is, but what it does in the academic context. A voluntary tool for reformulating notes differs from one used to support grading, flag misconduct, rank applicants or recommend progression. The same model may thus be low-impact in one setting and high-impact or potentially high-risk in another.

The regulatory position of an educational LLM therefore depends on its transition from general-purpose model to specific educational use case. At model level, the provider may be subject to GPAI obligations under the AI Act. At deployment level, the university must assess the system's intended purpose, workflow, data processing, human oversight and effects on academic rights, opportunities and outcomes. The key question is not simply whether an LLM is used, but whether and how its outputs influence educational processes or decisions affecting students, teachers or applicants.

2.2 Risk tiers and educational use cases

The AI Act does not classify LLMs used in higher education as high-risk by default. Its regulatory approach depends on the intended purpose of the AI system, the context in which it is deployed, and the possible impact on individuals. For the purposes of this section, four legal regimes are particularly relevant: a) prohibited AI practices; b) high-risk AI systems, including certain systems used in education and vocational training; c) AI systems subject to specific transparency obligations; d) and GPAI systems, which are regulated at model-provider level. Beyond these legal regimes, many educational uses of LLMs may remain low-impact from a legal perspective, while still requiring institutional guidance where they affect teaching, learning, academic integrity, student data or access to educational opportunities (European Union, 2024).

The typology below is therefore not presented as a formal legal taxonomy of the AI Act. It is an operational typology for higher education, developed by connecting the AI Act's risk-based logic with the ADMIT ethical dimensions and guidance on responsible GenAI use in academic settings (Carbonel et al., 2024; Triantafyllopoulos et al., 2025b). Its purpose is to help institutions distinguish among different families of use and identify the level of scrutiny, documentation and safeguards that each family may require.

Table 1 - Typology of LLM Use Cases in Higher Education: Risk Profile and Regulatory-Ethical Treatment

Use-case family	Typical examples in higher education	Main institutional risk	Indicative regulatory and ethical treatment
Optional individual study support	Brainstorming, summarisation, translation, language revision, coding assistance, self-testing, exploratory tutoring	Over-reliance, hallucinated content, plagiarism, undisclosed use, unequal access, unsafe data entry	Usually low-impact from a legal perspective, where it does not determine or materially influence institutional decisions; still requires guidance on academic integrity, disclosure, verification, data protection and responsible use
Teacher-mediated pedagogical support	Generation of examples, preparation of exercises, formative feedback drafts, quiz generation, teaching-material design, simulation prompts	Inaccuracy, bias, misalignment with learning outcomes, excessive delegation of pedagogical judgement	Requires teacher responsibility, quality control, accuracy checking, accessibility review and alignment with course objectives
Institutionally provided assistive tools	University-licensed chatbot, AI writing assistant, integrated learning-platform assistant, institutional support tool	Institutional endorsement, access inequality, data governance, procurement risks, unclear responsibility	Requires institutional disclosure, approved-use policies, accessibility guarantees, data protection review, user support and monitoring
Transparency-sensitive student interaction	AI helpdesk, tutoring bot, administrative chatbot, career guidance bot, orientation assistant, wellbeing triage tool	Users may misunderstand the nature, authority or reliability of AI-generated advice	Requires clear AI disclosure, explanation of limitations, escalation to human support, logging rules and safeguards for sensitive data
AI-assisted assessment and feedback	Automated or semi-automated feedback, rubric-based assessment support, AI-generated grading suggestions, essay evaluation support	Unfair or inaccurate evaluation, hidden influence on grades, lack of contestability, bias	Requires human oversight, calibration, documentation, student information, human review and appeal mechanisms; may require high-risk assessment where it materially influences academic outcomes
Admission, ranking and allocation	Applicant screening, admission recommendation, ranking of candidates, allocation to programmes or learning opportunities	Discrimination, exclusion, opacity, unequal access to educational opportunities	Potentially high-risk where it affects access, admission or allocation of educational opportunities; requires strict accountability, bias assessment, auditability, human review and procedural safeguards
Academic integrity and misconduct detection	AI detection, plagiarism analysis, cheating-risk scoring, behavioural	False positives, reputational harm, anxiety, surveillance, procedural unfairness	Potentially high-impact and possibly high-risk where linked to examination monitoring or disciplinary

	flagging, proctoring support		procedures; requires evidence access, human review, appeal and proportionality
Progression and learning pathway decisions	Recommendations on progression, remedial pathways, module access, risk-of-failure profiling, dropout-risk alerts	Profiling, self-fulfilling effects, unequal treatment, limitation of opportunities	Requires careful classification, transparency, human oversight, anti-discrimination safeguards and documentation, especially where decisions materially affect academic trajectories

Source: authors' elaboration based on the risk-based approach of Regulation (EU) 2024/1689, the ADMIT ethical taxonomy for Generative AI in higher education, and the ADMIT Guide of ethical use of LLM / Generative AI systems in Higher Education (European Union, 2024; Carbonel et al., 2024; Triantafyllopoulos et al., 2025b).

This typology shows that classification and governance should depend on the system's function, institutional integration, human involvement and consequences for students. The greater the impact on academic rights, opportunities or outcomes, the stronger the safeguards required. Optional, low-impact uses may be governed through lighter measures, provided institutions still address academic integrity, transparency, data protection and equitable access.

Classification should also distinguish advisory from determinative uses. A teacher-reviewed LLM suggestion has a different institutional effect from an automated recommendation that contributes to a grade, misconduct flag or admission decision. However, an output should not be treated as merely advisory simply because a human is formally involved: if the reviewer lacks the time, competence, authority or information to challenge it, the system may in practice become decisive.

Institutions should therefore document the intended purpose of each GenAI deployment, including whether it is optional or mandatory, who uses it, whether outputs are advisory or decision-shaping, what data are processed, whether formal academic outcomes are affected, and what oversight and appeal mechanisms exist. This supports both AI Act compliance and institutional accountability.

2.3 Applicability of Annex III to higher education

Annex III of the AI Act is especially relevant to higher education because it includes AI systems used in education and vocational training to determine access or admission, assign persons to institutions, evaluate learning outcomes, assess educational level or monitor prohibited behaviour during tests (European Union, 2024). Its relevance depends not on the use of an LLM as such, but on whether the system materially influences access, assessment, progression, allocation of opportunities or examination monitoring. In universities, this may include LLM-based support for application summaries, ranking, admission recommendations, grading, assessment feedback, misconduct flags, proctoring, remedial pathways or progression decisions. These uses are sensitive

because they may affect students’ opportunities, reputation, certification, mobility, scholarships or future trajectories.

Annex III should not be applied mechanically. Under Article 6, an Annex III system may fall outside the high-risk category where it does not pose a significant risk to health, safety or fundamental rights, including where it does not materially influence decision-making. This is important for universities, since many LLM tools perform preparatory, procedural or assistive tasks without directly shaping academic decisions. However, the “non-material influence” qualification should be applied cautiously: a system formally described as advisory may become decisive if staff rely heavily on AI-generated summaries, scores, rankings, indicators or alerts, or if AI outputs determine which cases are prioritised, filtered, escalated or made visible. The key question is whether human reviewers have the time, competence, authority and information to critically assess, disregard or override the AI output.

Profiling also requires attention. Where an Annex III system profiles natural persons, the AI Act treats this as particularly sensitive. In higher education, learning analytics, dropout-risk prediction, academic performance profiling, behavioural monitoring and adaptive pathway systems may shape how students are perceived, supported, monitored or guided, and may create self-fulfilling effects when “at risk” students receive different expectations, opportunities or treatment.

Biometric and emotion-recognition technologies require separate classification assessment. In higher education, they may appear in AI proctoring, identity verification, classroom monitoring, behavioural analysis or engagement detection. If an LLM-based environment is combined with biometric categorisation, remote biometric identification, emotion recognition or behavioural monitoring, institutions should assess both Annex III education provisions and wider fundamental-rights implications, especially in mandatory or high-stakes settings.

The key policy implication is that universities need a documented classification process for institutional GenAI deployments. Before adopting an LLM-based system in formal educational or administrative processes, institutions should identify its purpose, users, affected persons, data processed, role in the workflow, potential impact on academic decisions and safeguards. This helps avoid under-classifying high-impact systems as ordinary tools, while also avoiding over-classifying optional, low-impact pedagogical uses as high-risk.

2.4 General-purpose AI obligations and the higher education deployment context

Articles 51–55 of the AI Act establish a specific regime for GPAI models, with additional obligations for models posing systemic risk (European Union, 2024). Although addressed mainly to model providers rather than universities, these provisions are relevant because many institutional GenAI deployments rely on third-party GPAI models integrated into learning platforms, tutoring systems, assessment tools, administrative workflows or student-support services. The distinction between model provider and institutional deployer is therefore central. Providers develop or make GPAI models available; universities usually procure, license, configure, integrate or recommend LLM-

based tools for specific educational or administrative purposes. In doing so, universities do not assume responsibility for model development, but they are responsible for the institutional conditions of use, including tool selection, deployment purpose, user groups, data processing, human oversight, communication to students and staff, and safeguards for academic decision-making.

The AI Act requires GPAI providers to maintain technical documentation on training, testing and evaluation, provide relevant information to downstream actors, establish a copyright compliance policy and publish a sufficiently detailed summary of training content. For universities, this information is essential for assessing whether a model is suitable for a given educational use. That assessment should cover not only technical performance, but also pedagogical suitability, transparency, hallucination and bias risks, accessibility, data protection, cybersecurity, intellectual property, contractual terms, model updates and impacts on students and staff. Provider compliance is therefore necessary but insufficient: a documented GPAI model may still be deployed in ways that are pedagogically inappropriate, opaque, unfair or unsafe, especially in grading, proctoring, academic integrity investigations or student-risk profiling.

Procurement is a key link between the GPAI regime and higher education governance. When procuring or licensing LLM-based systems, universities should require information on the model or service used, model version, update practices, known limitations, data processing and retention, security, accessibility, incident notification, subcontractors, and whether user inputs may be stored, reused or transferred. They should also verify whether the system relies on a GPAI model with systemic risk and, where relevant, whether the provider has met the related obligations.

Model versioning also requires attention. LLM-based systems may change through updates, fine-tuning, retrieval augmentation, safety adjustments or provider policy changes, affecting consistency, reliability and fairness. Tools used for feedback, assessment support or student guidance should not change behaviour invisibly. Universities should therefore require update notifications and procedures for reassessing educational suitability after significant model or system changes.

The higher education context requires a layered interpretation of GPAI obligations. Providers must meet AI Act model-level requirements; vendors and integrators must provide documentation and contractual guarantees; universities must decide whether, where and how systems are used; teachers, students and staff need practical guidance; and governance and quality assurance bodies must monitor impacts, incidents, complaints and improvement needs. This avoids both over-reliance on provider compliance and treating universities as responsible for technical model-development choices beyond their control. Provider obligations and institutional responsibilities are distinct but interdependent. EU policy should therefore clarify how GPAI provider obligations can be operationalised in higher education procurement, deployment and governance.

2.5 Cross-reference with the ADMIT ethical taxonomy

The AI Act provides the legal baseline for classifying AI systems and identifying the obligations associated with different regulatory regimes. It is essential for determining whether a given LLM-based system falls within prohibited practices, high-risk AI systems, transparency obligations, general-purpose AI model obligations or lower-impact uses. However, legal classification alone does not fully capture the educational, pedagogical and institutional implications of LLM deployment in higher education. For this reason, the ADMIT ethical taxonomy provides a complementary framework for translating legal classification into concrete governance questions for universities (Carbonel et al., 2024).

The ADMIT taxonomy covers eight ethical dimensions: transparency, accountability, technical robustness and safety, educational impact and academic integrity, privacy and data governance, societal, individual and environmental wellbeing, teacher and student agency and oversight, and diversity, non-discrimination and fairness. All these dimensions are relevant to responsible GenAI governance in higher education and none should be considered as less important than any other. This section therefore provides a cross-reference between all eight ADMIT dimensions and the AI Act, showing how each connects with the operational conditions that institutions must establish when classifying, procuring, deploying and monitoring LLM-based systems. The ADMIT project does not rank these dimensions hierarchically: each addresses a distinct and essential aspect of responsible GenAI governance that applies across the three main HE stakeholder perspectives identified in the ADMIT project - students, teachers and institutions (including administrative and support layers).

The following table provides a crosswalk between all eight ADMIT ethical dimensions, the classification logic of the AI Act and the institutional safeguards that universities may need to adopt. Its purpose is to show how each dimension helps translate AI Act classification into practical governance requirements for higher education institutions, addressing students, teachers and institutional leadership as the three key stakeholder perspectives.

Table 2. Crosswalk between all eight ADMIT Ethical Dimensions, AI Act Classification and Institutional Safeguards

ADMIT dimension	Core institutional question	Relevance for AI Act classification	Examples of safeguards in higher education
Transparency	Do students, educators and institutional stakeholders understand which AI system is used, for what purpose, with what limits and with what possible consequences?	Supports informed use, user awareness, traceability and meaningful communication about AI-mediated processes, especially where systems interact with users, generate content, support assessment or	AI-use notices, student-facing explanations, disclosure of the AI role in assessment or support services, model documentation summaries, explanation of limitations, information on human review, escalation channels and contestability procedures

		influence academic decisions	
Accountability	Are responsibilities clearly assigned for procurement, deployment, monitoring, complaints, review and redress?	Connects provider and deployer obligations with internal university governance, especially where institutional use of AI may affect students, applicants, educators or academic decision-making processes	Governance roles, approval procedures, procurement and service-level clauses, complaint routes, appeal mechanisms, incident reporting, documented human oversight, periodic review of institutional deployments
Technical Robustness and Safety	Is the system reliable, secure, regularly checked, resilient to misuse and supported by contingency plans?	Supports risk management, accuracy, robustness, cybersecurity and continuity of academic processes, particularly where AI tools are embedded in formal teaching, assessment, student support or administrative workflows	Security checks, data breach procedures, prompt-injection mitigation, output validation, update monitoring, fallback procedures, continuity plans, monitoring of errors, vulnerabilities and model changes
Educational Impact and Academic Integrity	Is GenAI use pedagogically justified, aligned with learning outcomes and compatible with authentic assessment and academic integrity?	Connects with high-risk classification where AI affects assessment validity, examination monitoring or academic decision-making; supports proportionate governance of integrity-sensitive uses	Clear rules on permitted, restricted and prohibited GenAI uses; disclosure norms; assessment redesign; fair misconduct procedures; human review of AI-generated suspicion; alignment of AI use with intended learning outcomes
Privacy and Data Governance	Are student and institutional data protected, processed lawfully and not retained or reused beyond the defined purpose?	Supports GDPR compliance, identification of an appropriate lawful basis, data minimisation and, where consent is relied upon, valid and freely given consent; connects with AI Act documentation and GPAI transparency requirements for data processing	Data protection impact assessments, procurement clauses on retention and reuse, restrictions on uploading student work to external systems, anonymisation policies, vendor data-handling agreements, student information on data processing
Societal, Individual and	Does the deployment avoid over-reliance, stress,	Supports proportionality assessment and attention to broader	Monitoring of student stress and over-reliance, assessment of

Environmental Wellbeing	social harm and disproportionate environmental costs?	societal and environmental impacts, particularly where mandatory or large-scale institutional deployment is involved	environmental costs, guidance on responsible use, institutional review of social and psychological impacts, proportionality checks for mandatory AI-mediated processes
Teacher and Student Agency and Oversight	Do educators and students have adequate AI literacy, human control, choice in using GenAI and fair alternatives?	Supports the AI Act's human oversight requirements and AI literacy provisions; ensures that GenAI empowers rather than replaces human judgement in academic contexts	Structured AI literacy programmes for students and staff, teacher training on AI verification and assessment redesign, opt-out mechanisms, freedom to choose non-AI alternatives, meaningful human oversight in high-stakes decisions
Diversity, Non-discrimination and Fairness	Are GenAI tools accessible, free from discriminatory bias and equitably available regardless of disability, location, language or socio-economic background?	Connects with the AI Act's non-discrimination and fundamental rights provisions; supports equitable treatment in high-risk and Annex III educational use cases	Accessibility-by-design, compatibility with assistive technologies, bias monitoring, equitable institutional access, non-disadvantage principle for opt-out, monitoring of unequal effects on specific student groups, multilingual support where relevant

Source: authors' elaboration based on Regulation (EU) 2024/1689, the ADMIT ethical taxonomy and the ADMIT Institutional Self-Assessment Tool (European Union, 2024; Carbonel et al., 2024; Triantafyllopoulos et al., 2025a).

This cross-reference shows that the AI Act should not be used as a stand-alone checklist. While it identifies legal status and obligations, the eight ADMIT ethical dimensions define the full set of institutional conditions that make compliance meaningful in academic practice. Together, they support a governance model that is legally aware, institutionally accountable and operationally implementable: users must understand when and how AI is used (transparency); responsibilities must be clearly assigned across the full deployment chain (accountability); systems must be reliable, secure and controllable in real educational settings (technical robustness and safety); pedagogical justification and alignment with learning outcomes must be ensured (educational impact and academic integrity); student and institutional information must be protected (privacy and data governance); over-reliance, stress, social impact and environmental costs must be addressed (wellbeing); AI literacy, human control and fair alternatives must be guaranteed (teacher and student agency and oversight); and accessibility, bias monitoring and safeguards against unequal treatment must be in place (diversity, non-discrimination and fairness). The ADMIT taxonomy thus provides a comprehensive framework for augmenting the AI Act's legal obligations with sector-specific ethical dimensions tailored to higher education, paving the way for future sectoral contextualisation of the

AI Act. This approach also aligns with the ADMIT Institutional Self-Assessment Tool, which turns the taxonomy into role-specific questions for students, educators and institutional leadership (Triantafyllopoulos et al., 2025a). Questions on understandability and contestability, responsibility for AI-supported teaching or assessment, and procurement, monitoring and technical continuity all point to the same governance need: GenAI deployment must be transparent, accountable and technically robust in actual institutional practice.

2.6 Implications for policy recommendations

The classification of educational LLMs under the AI Act therefore has an important but limited function. It helps higher education institutions identify when an LLM-based application remains an assistive or transparency-sensitive tool and when it becomes part of a process that may affect access, assessment, progression or other significant academic outcomes. However, classification alone does not solve the institutional governance problem. Once a university has identified the relevant risk level, it still needs to decide how to organise transparency, accountability, human oversight, academic integrity, data governance, procurement, accessibility, teacher training and quality assurance. This is where a sector-specific policy gap emerges. The AI Act provides the legal architecture for regulating AI systems and GPAI models, but it does not define in detail how higher education institutions should translate these obligations into academic rules, pedagogical practices, institutional responsibilities and student-facing safeguards.

The classification of LLM-based systems in higher education should therefore be linked to proportionate institutional safeguards. When a system affects admission, assessment, progression, monitoring or academic integrity, the institution should treat human oversight, contestability, documentation, bias monitoring, accessibility, data protection and appeal procedures as essential requirements. When a system is used only for optional support, lighter obligations may be appropriate, but institutions should still ensure transparency, academic integrity, equity of access and responsible use. The AI Act should not be interpreted as a complete governance framework for LLMs in higher education. It provides the legal architecture for risk classification, GPAI obligations and high-risk use cases, but the specific translation of these principles into academic policies, quality assurance procedures, teacher responsibilities, student rights and institutional monitoring remains a sectoral policy task - one to which the present deliverable, D5.5 of the ADMIT project, directly contributes by proposing concrete recommendations grounded in the ADMIT ethical taxonomy and operational tools (D5.2, D5.3, D5.4). In this sense, D5.5 serves as the basis for sector-specific EU guidance that interprets the AI Act from a higher education perspective and augments its legal obligations with the less formal but equally essential ethical dimensions identified by the ADMIT project. The following subsection therefore moves from classification to gap analysis. It identifies the areas in which the current EU regulatory and policy framework still requires operational translation for the specific context of higher education.

3 Current Regulatory, Policy and Implementation Gaps in Relation to Higher Education

The analysis of educational LLMs under the AI Act shows that European regulation already provides a strong legal baseline. The Act adopts a risk-based approach, sets obligations for providers and deployers, regulates GPAI models, and identifies high-risk educational uses where AI affects access, admission, assessment, allocation, progression or examination monitoring (European Union, 2024). These provisions are directly relevant to LLMs and Generative AI in universities. However, higher education raises issues that legal classification alone cannot resolve. Universities embed technologies in pedagogical relationships, assessment cultures, student support, academic integrity, quality assurance and institutional governance. The same AI system may therefore create different risks depending on whether it is used informally, recommended by a teacher, embedded in a learning platform, integrated into assessment or linked to administrative decision-making. As defined in the ADMIT project Description of the Action, the analysis of these risks should be structured around the three key stakeholder perspectives: students, who need to develop digital skills, critical thinking and ethical awareness for responsible GenAI use; teachers, who are responsible for designing learning activities, verifying AI outputs and ensuring academic integrity; and institutions (including leadership, administrative and support layers), which must establish governance frameworks, procurement standards, quality assurance procedures and the necessary educational and ethical conditions for responsible GenAI deployment.

The practical challenge lies in the “last mile” of governance: how broad European principles are converted into everyday institutional decisions. Universities need criteria for deciding when GenAI use requires formal approval, what evidence must be collected before deployment, how responsibilities are allocated across academic, technical and administrative units, and how safeguards are made visible to students and staff. Without this operational layer, compliance may remain formally correct but uneven in practice, with similar tools governed differently across programmes, institutions and national contexts.

3.1 The AI Act does not define a sector-specific governance model for universities

The AI Act regulates AI systems through roles such as provider, deployer, importer and distributor. This horizontal structure is necessary, but it does not map neatly onto the internal organisation of universities, where leadership, faculties, departments, teachers, students, IT services, data protection officers, quality assurance units and vendors all shape digital deployment. The Act does not specify who should approve GenAI tools, what review is needed before deployment, how teachers and students should be supported and informed, how complaints should be handled, or how policies should be updated. A university may therefore meet general deployer obligations while still lacking a coherent internal framework for deciding when GenAI may be used, under what conditions, with which safeguards and accountability mechanisms. This gap is especially

relevant where deployment emerges through fragmented experiments, local initiatives, vendor-led solutions or informal individual use.

A sector-specific governance model should connect strategic responsibility, procurement, academic policy, teacher practice, student rights, data governance and quality assurance. It should define responsibility for tool approval, acceptable uses, risk review, deployment monitoring, user communication, incident handling and policy updates. These functions cannot be left to providers, which do not control the academic context of use. EU-level guidance should therefore translate AI Act obligations into university governance practices, preserving institutional autonomy while setting a shared baseline for approval, risk-based review, documentation, stakeholder involvement, human oversight, complaints and quality assurance integration. The present deliverable aims to contribute to this translation by providing considerations, guidelines and examples for each of the stakeholders in higher education - students, teachers and institutions - along the dimensions of the ADMIT ethical taxonomy, as mandated in the ADMIT Description of the Action.

3.2 Academic integrity is not addressed as a distinct regulatory category

A major gap concerns academic integrity, identified by ADMIT as a core dimension of responsible GenAI use in higher education (Triantafyllopoulos et al., 2025b). While the AI Act addresses risks to health, safety and fundamental rights, and classifies some educational uses as high-risk when they affect access, assessment, progression or examination monitoring, it does not treat academic integrity as a distinct regulatory category. This is a critical limitation, since academic integrity underpins trust in learning, assessment, certification and research.

LLMs can support legitimate academic work by helping students brainstorm, improve language, test understanding, receive formative explanations or obtain coding assistance, and by helping educators design materials, exercises, feedback drafts and assessment formats. These uses may be valuable when transparent, proportionate and aligned with learning outcomes. At the same time, LLMs can blur the line between legitimate support and misrepresentation of authorship, enabling hidden outsourcing of cognitive work, fabricated references, unverified claims, inappropriate paraphrasing, ghost-writing, code without comprehension and submission of AI-generated work as original. The risks therefore concern not only plagiarism, but assessment validity, authentic performance, critical thinking and the credibility of qualifications. Academic integrity should therefore not be reduced to misconduct detection. A response based mainly on AI-detection tools would be inadequate and potentially harmful, given risks of unreliability, false positives, disadvantage for non-native speakers or particular writing styles, and adversarial relations between students and institutions. A better approach treats academic integrity as a design and governance issue. Institutions should define clear, understandable and discipline-sensitive rules on permitted, restricted and prohibited LLM uses in coursework, examinations, research-related tasks and student support. For example, language refinement may be acceptable in some contexts but inappropriate where writing ability is assessed; coding assistance may be permitted only if students can explain, adapt and justify the code.

Disclosure is central to this framework. Students should know when GenAI use is allowed or required to be declared, what declaration is expected, and how it will be evaluated. Disclosure should support transparency, reflection and responsible learning rather than operate only as a punitive tool. Documenting prompts, accepted or rejected outputs and verification steps may provide better evidence of learning than post-submission AI detection.

Assessment redesign is equally important. Where traditional take-home assignments no longer provide reliable evidence of individual learning, institutions should support formats such as oral components, in-person or online supervised writing, process portfolios, reflective commentaries, supervised stages, viva-style discussion, authentic tasks, project logs, iterative drafts, peer review, contextualised problem-solving and critical evaluation of AI outputs. This should be seen not only as a defensive response to GenAI, but as an opportunity to improve assessment validity and alignment with learning outcomes. The appropriate mix of formats depends on the institution's legal framework and delivery model. Distance higher education institutions, whose students learn entirely or predominantly online and whose assessment arrangements are governed by specific national regulations, may combine online AI-proctored examinations, synchronous oral examinations via videoconference, and in-person supervised sessions where required by their national regulatory framework. EU guidance should treat all these arrangements as legitimate alternatives, rather than privileging any single format as the default standard.

Academic integrity policies should also protect students from unfair or automated accusations. AI-generated suspicion should never be treated as proof of misconduct. Flags from detection tools, proctoring systems or automated analyses should trigger competent human review, access to relevant evidence where appropriate, the opportunity to respond, and a fair and proportionate procedure. Students should not have to disprove opaque algorithmic suspicion without meaningful explanation.

The current EU framework provides no minimum standards for these issues. It does not define how institutions should distinguish acceptable assistance from substitution of student work, establish disclosure norms, safeguard AI-detection or misconduct procedures, require assessment redesign, or connect academic integrity with student rights, teacher responsibilities and quality assurance. This creates a risk of fragmentation across the European Higher Education Area. EU guidance could provide a common baseline while preserving disciplinary and institutional autonomy: clear rules on permitted and prohibited uses, disclosure expectations, assessment redesign principles, safeguards against unreliable detection, human review of suspected misconduct, proportional procedures, student support and integration into internal quality assurance.

3.3 Existing EU frameworks do not provide sufficient transparency requirements for educational GenAI use

Transparency is central to legal and ethical AI governance, but the current EU framework does not yet provide sufficiently detailed standards for LLM and GenAI deployment in higher

education (European Union, 2024; Carbonel et al., 2024). The AI Act includes transparency provisions and documentation duties for GPAI providers, but universities need more concrete, context-sensitive standards for communicating AI use to students, educators, staff and other affected stakeholders.

In universities, transparency cannot be reduced to a generic statement that “AI is used”. Meaningful disclosure should explain which system is used, for what purpose, at which stage of the academic workflow, with what data, limits, human involvement and consequences for learning, assessment, support or decision-making. Because LLMs may support teaching materials, feedback, tutoring, administration, grading, misconduct detection, proctoring, student-record summaries, interventions, admission or progression decisions, disclosure should be proportionate to the use case. Optional study support may require a visible AI notice and responsible-use guidance, while assessment, proctoring or academic integrity uses require more detailed information on the system’s role, safeguards and rights to human review.

The EU framework does not define a common format for institutional GenAI disclosure in higher education. It does not require registers of AI systems used in teaching, learning, assessment, student support or administration, nor minimum content for student-facing notices or rules on when students must be informed that AI contributed to feedback, evaluation, proctoring, profiling or guidance. Nor does it clarify how GPAI provider information should be translated into accessible explanations for educational users. This creates a gap between provider-level and institutional transparency: technical documentation may be too complex, general or vendor-oriented for students, teachers or governance bodies, so universities need to convert it into understandable and actionable disclosure.

A student-facing disclosure standard should state the AI system or type used, its purpose, the academic process involved, the nature of its output, whether it is advisory or decision-shaping, what personal or assessment-related data are processed, whether data may be stored or reused, known limitations and risks of inaccurate or biased outputs, and how students can request clarification, human review or redress. Educators need additional operational information on approved uses, permitted data inputs, verification duties, disclosure requirements and procedures for problematic outputs. Institutions need documentation for procurement, monitoring, auditability and policy review.

Transparency is also linked to academic integrity, fairness and inclusion. Students cannot follow responsible-use rules unless institutions clearly explain when GenAI is permitted, restricted, prohibited or must be declared. Clear communication also reduces asymmetries between students with different digital competences. In higher education, transparency should therefore function as an institutional architecture, including approved-system inventories, procurement documentation, staff guidance, student notices, course-level rules, assessment instructions, appeal information, data protection information and quality assurance review. The level of detail should be proportionate to risk, but affected persons should always understand AI’s role and available safeguards.

The gap is that no EU-level instrument yet provides this sector-specific transparency architecture for LLMs in higher education. The AI Act offers legal and technical foundations, especially for GPAI documentation and high-risk systems, while the ADMIT taxonomy and tools provide an ethical and operational framework. Universities still need practical policy guidance to turn these foundations into disclosure practices that are meaningful for students, usable for educators and manageable for institutions.

3.4 The provider-deployer distinction leaves unresolved accountability questions

The AI Act's provider–deployer distinction is essential, but it leaves unresolved accountability questions in higher education. Universities often rely on third-party LLMs, platforms, proctoring services, tutoring systems, assessment tools or administrative applications. Providers control the model, documentation, updates and security, while universities control educational purpose, institutional integration, academic workflows and communication with students and staff. This creates a potential accountability gap. If a system produces misleading feedback, biased recommendations, unfair grading suggestions, inaccurate tutoring, harmful profiling, false misconduct flags or inappropriate academic advice, responsibility cannot rest only with the provider or an individual teacher. Providers are responsible within their technical and contractual perimeter; universities are responsible for adoption, purpose, conditions of use, safeguards, monitoring, review and redress.

Accountability is therefore not only legal liability after harm, but a governance question before deployment. Institutions should be able to identify who proposed the system, reviewed its educational suitability, assessed data protection and cybersecurity, approved procurement, defined acceptable use, trained staff, informed students, monitored performance and handled complaints or incidents. Without this responsibility chain, GenAI deployment risks becoming fragmented, informal or dependent on individual initiative. Accountability must also reflect academic harms, which may include unfair assessment, exclusion from opportunities, unjustified misconduct suspicion, reputational damage, discriminatory treatment, psychological stress, loss of trust, inappropriate disclosure of student data, or reduced learning quality. These harms may be difficult to quantify but can be highly consequential for students and institutional legitimacy. A clear accountability model should distinguish five linked levels: provider accountability, covering the technical system, documentation, updates, cybersecurity, contracts and AI Act obligations; procurement accountability, covering the university's decision to acquire or approve a system after reviewing suitability, risks, conditions and institutional values; deployment accountability, covering configuration, integration, communication and monitoring in specific processes; pedagogical accountability, covering teachers' use of the system in relation to learning outcomes, assessment, feedback, academic integrity and student support; and procedural accountability, covering complaints, explanations, challenges, human review and redress. These levels form a chain: weaknesses in documentation, deployment rules, pedagogical responsibility, complaints or quality assurance can make AI-related problems invisible or difficult to contest.

Accountability should also address the position of academic, administrative and support staff as workers affected by institutional AI deployment. Under Article 26(7) of the AI Act, employers deploying high-risk AI systems in the workplace must inform affected workers and their representatives before deployment. Beyond this formal requirement, governance arrangements should protect academic freedom and professional autonomy, assess effects on workload and role boundaries, and prohibit the use of opaque AI-generated indicators as the sole basis for staff evaluation or managerial decisions.

Service-level agreements and contracts are important but insufficient. For institutional GenAI systems, they should define support, maintenance, incident notification, updates, data processing, cybersecurity, documentation, accessibility, continuity and termination. Yet contractual safeguards do not replace student transparency, teacher guidance or institutional oversight.

The ADMIT Institutional Self-Assessment Tool supports this approach by operationalising accountability through questions on service-level arrangements, institutional responsibilities, complaints and redress (Triantafyllopoulos et al., 2025a). It asks whether support and maintenance responsibilities are defined, whether responsibilities for GenAI-related problems are assigned, and whether educators and students can report concerns or request redress through accessible procedures. These questions translate accountability into concrete institutional requirements. The current gap is that no EU-level instrument defines this accountability architecture for GenAI in higher education. The AI Act provides the provider–deployer structure, ADMIT identifies the institutional questions, and the ESG offer the quality assurance logic. What is missing is sector-specific EU guidance connecting these elements into a practical university governance model.

3.5 Human oversight and appeal mechanisms are not specified for academic contexts

The AI Act recognises human oversight, especially for high-risk systems (European Union, 2024), but higher education needs a more precise definition of meaningful oversight in academic practice. A human in the loop is not enough if the reviewer lacks the time, competence, authority, information or mandate to assess and, where needed, override AI outputs. In universities, systems labelled as advisory may still shape grading, misconduct flags, admission rankings, student-risk profiles or proctoring alerts. Oversight becomes meaningful only when reviewers can understand the system’s purpose, output, limits, data basis, possible errors or bias, and consequences, and have the authority to disregard, correct or escalate the result.

In assessment, AI-generated feedback or grading suggestions should not be treated as neutral evidence. Teachers remain responsible for checking alignment with learning outcomes, criteria, disciplinary standards and the student work submitted. Institutions should define how systems are calibrated, outputs checked, disagreements handled and clarification or review requested.

In admission and selection, oversight should prevent AI-generated rankings, summaries or recommendations from obscuring candidate evaluation. Committees should be able to inspect the

basis of AI-supported recommendations, identify bias, consider context and avoid reducing admission to an opaque scoring process.

In academic integrity and misconduct procedures, oversight is particularly sensitive. AI detection, plagiarism tools, behavioural monitoring and proctoring systems may produce false positives or misleading indicators. AI flags should never be treated as proof of misconduct, but only as signals requiring contextual human review, procedural fairness, access to relevant evidence where appropriate, and the opportunity for students to respond.

AI-proctored examinations illustrate this risk. Alerts about gaze direction, background noise, face absence, typing patterns, connection problems or unusual movements may reflect disability, anxiety, technical conditions, cultural behaviour, domestic environments or harmless interruptions. Human review should consider this context before any academic consequence is imposed. Online proctoring is of particular relevance for distance higher education institutions, where remote examination is part of the standard assessment model - whether as the primary mode or alongside in-person supervised sessions. In these settings, the volume and frequency of AI-proctored interactions is substantially higher than in campus-based institutions, and the consequences of miscalibrated alerts, false positives, or technically generated anomalies are correspondingly more significant. EU guidance should therefore require that AI proctoring systems used in any higher education context be validated for the specific examination environment in which they are deployed, including remote domestic settings, and that human review procedures be proportionate to the scale and frequency of proctoring use. Appeal mechanisms are the procedural counterpart of oversight: students should be able to challenge AI-mediated outcomes in assessment, grading, admission, proctoring, misconduct procedures, progression decisions or allocation of opportunities. Effective appeals require notice that AI contributed to the decision or flag, an explanation of its role, access to relevant evidence subject to legitimate limits, review by a competent human decision-maker, a reasoned outcome and documentation for accountability and quality improvement. The current EU framework does not specify these safeguards for academic contexts.

Human oversight also depends on AI literacy and staff training. Teachers, exam boards, admission committees and administrative staff need to understand system limits, automation bias, hallucination, false positives, data protection, evidence evaluation, bias and student rights. Oversight is therefore both an organisational and a competence requirement.

The ADMIT dimension of Teacher and Student Agency and Oversight is especially relevant here (Carbonel et al., 2024). It stresses that GenAI should empower users, preserve human judgement, allow intervention where contextual understanding is needed, and avoid misleading or manipulating students. In higher education, oversight should therefore preserve the human and pedagogical nature of academic decision-making. The European Commission should support minimum standards clarifying that significant academic decisions must not rely solely on AI outputs; AI flags and recommendations require competent human review; affected students need meaningful information; appeals must be accessible and timely; and review outcomes should be documented for accountability and improvement.

3.6 The Digital Education Action Plan does not yet address the specific governance challenges of LLMs

The Digital Education Action Plan 2021–2027 provides an essential strategic framework for digital education in Europe (European Commission, 2020). It promotes a common vision of high-quality, inclusive and accessible digital education and supports the adaptation of education and training systems to the digital age. Its priorities concerning a high-performing digital education ecosystem, digital skills and competences, digital infrastructure, teacher support, AI-related competences, digital literacy and ethical guidance remain highly relevant for higher education. However, the Action Plan was not designed around the current scale, capabilities and institutional penetration of LLM-based Generative AI systems. The widespread availability of conversational AI, AI writing assistants, coding assistants, AI tutors, automated feedback tools, synthetic content generators and GenAI-enabled assessment support has created governance questions that go beyond the general digital transformation agenda. These questions concern authorship, academic integrity, assessment validity, transparency of AI-mediated feedback, student data entered into AI systems, procurement of third-party LLM tools, accessibility of paid AI services, teacher preparedness, and the role of human judgement in academic decisions.

The Digital Education Action Plan provides an enabling policy environment, but it does not define an operational framework for institutional GenAI governance in universities. It does not specify how higher education institutions should disclose the use of LLMs in teaching, assessment or student support. It does not define minimum safeguards for AI-assisted grading, proctoring, academic integrity checks or student-risk profiling. It does not establish common expectations for opt-out mechanisms, appeal procedures, human review or the use of student data in AI interactions. It also does not provide procurement-oriented guidance for universities adopting commercial or open-source LLM-based systems. This gap is understandable. The Action Plan emerged before the large-scale public deployment of the current generation of LLMs and before GenAI became a daily tool for students, educators and institutions. At the time, the policy emphasis was appropriately placed on connectivity, digital equipment, digital skills, digital content, teacher guidance, interoperability, data use and the broader conditions for digital transformation. These priorities remain necessary, but they are no longer sufficient to address the specific risks and opportunities introduced by GenAI in higher education.

The current generation of LLMs differs from many earlier digital education technologies because it generates academic content, interacts conversationally with users, can be integrated into many institutional processes, and may process prompts, student work, personal information or institutional data in ways that are not always transparent. For this reason, LLM governance cannot be treated only as a subset of general digital competence. Students and educators need to understand how LLMs generate outputs, why they may hallucinate, how bias may appear, when AI assistance must be disclosed, how generated information should be verified, what data should not be entered into external systems, and how AI use interacts with learning outcomes and academic integrity.

The Action Plan also does not fully address the new relationship between institutional provision and private access to AI tools. If some students have access to more powerful commercial LLMs while others rely on free, limited or institutionally unavailable tools, GenAI may create a new layer of inequality. This is not only a technical access issue. It affects the fairness of learning opportunities, the validity of assessment, and the capacity of students to participate in AI-enhanced learning activities.

From the perspective of ADMIT, the Digital Education Action Plan should therefore be seen as a strategic platform that needs a GenAI-specific extension for higher education. The ADMIT taxonomy, guide and institutional self-assessment tool can contribute to this extension by providing ethical dimensions, role-specific guidance and operational monitoring instruments. These resources can help translate the Action Plan's broad objectives into concrete institutional practices concerning transparency, accountability, technical robustness, academic integrity, agency, fairness, privacy and wellbeing.

From a policy perspective, the European Commission should update or complement the Digital Education Action Plan with guidance specifically addressing LLMs and Generative AI in higher education. This guidance should connect AI literacy with academic integrity, inclusion with equitable access to GenAI tools, digital transformation with institutional accountability, and innovation with safeguards for students and educators. In this way, the Action Plan could continue to provide the strategic umbrella for digital education while becoming more responsive to the concrete governance challenges introduced by LLMs.

3.7 Equity, accessibility and opt-out mechanisms remain insufficiently regulated

The ethical use of LLMs and Generative AI in higher education raises equity, accessibility and student-agency issues that current EU frameworks do not yet address operationally. The Digital Education Action Plan promotes inclusive and accessible digital education, and the ADMIT taxonomy covers diversity, non-discrimination, fairness, agency and oversight (European Commission, 2020; Carbonel et al., 2024). However, no EU-level framework specifies how universities should ensure equitable GenAI access, accessibility for students with disabilities, protection against AI-driven inequalities, or opt-out safeguards in high-stakes contexts. This gap matters because private access to stronger commercial models, premium subscriptions, advanced tools or better infrastructure may advantage some students in drafting, feedback, coding, translation, exam preparation, research support and confidence, amplifying existing inequalities if teaching and assessment assume comparable GenAI access.

The problem is not only financial. Access also depends on language, disability, digital literacy, location, connectivity, devices, interface design, cultural familiarity and confidence with conversational technologies. Students with disabilities may face barriers where tools are incompatible with assistive technologies, inaccessible, unavailable in usable formats or lacking alternative interaction modes. International and multilingual students may benefit from language

support, but may also face linguistic bias or uncertainty about permitted assistance. When a university integrates, recommends or designs learning activities around an LLM, it assumes responsibility for fair, accessible and pedagogically appropriate access. Formal academic activities should not advantage students who can buy stronger tools or already have higher AI literacy. Institutions therefore need transparent policies defining when GenAI is allowed, supported, required or prohibited, together with equivalent learning opportunities for students who do not use a specific tool.

Accessibility should be a design requirement. Before adopting GenAI in formal learning or support processes, institutions should assess usability for students with disabilities, compatibility with assistive technologies, language availability where relevant, clarity for users with different digital competences, and the availability of guidance and human support channels.

Bias and discrimination are closely linked to access. LLMs may reflect cultural, linguistic, gender, disability-related or socio-economic biases that affect examples, feedback, language evaluation, argument assessment, pathway recommendations or risk profiling. If used in assessment, student support or progression monitoring, such bias may become institutionally consequential. Universities therefore need procedures to detect, investigate and address unequal effects from GenAI deployment.

Linguistic performance should be addressed as a specific governance and quality issue, given the multilingual nature of European higher education. Institutions should evaluate GenAI performance in the languages and language varieties in which teaching, assessment and student support are provided, including lower-resource languages and non-native academic writing. Known language-specific limitations should be communicated to users, and systems should not be adopted for language-sensitive assessment or decision-making without evidence of adequate validity and fairness in the relevant linguistic context.

Opt-out is equally important. Students may have legitimate reasons to avoid a specific AI-mediated process, including privacy concerns, accessibility barriers, ethical objections, surveillance concerns, lack of trust or fear that AI may misinterpret their behaviour or work. In high-stakes contexts such as assessment, proctoring, academic integrity investigations or student-risk profiling, opt-out should function as a safeguard where GenAI may affect student rights, assessment validity, privacy, dignity, accessibility or academic opportunities. Where appropriate, alternatives should be educationally equivalent and should not create academic disadvantage.

The non-disadvantage principle is central: alternatives should preserve academic standards without penalising students for using a legitimate safeguard. This may require human-supervised assessment, oral examination, supervised in-person or online sessions, alternative submission routes, non-AI-mediated support or equivalent human review. The feasibility of specific alternatives depends on the institutional delivery model and the applicable national regulatory framework. Distance higher education institutions may not be able to offer in-person alternatives for all students, but they can provide functionally equivalent safeguards - such as synchronous oral examination via videoconference, human-reviewed online assessment, or alternative digital submission routes -

provided these are educationally equivalent and do not create academic disadvantage. EU guidance should avoid prescribing a single opt-out format and instead require that institutions demonstrate the availability of a meaningful, non-penalising alternative appropriate to their context. Equity and accessibility also require AI literacy support. Access to tools is insufficient if students cannot use them critically, safely and responsibly, verify outputs, avoid hallucinations, protect data, disclose use and challenge AI-mediated decisions.

The ADMIT self-assessment approach offers an operational basis for this gap. Its agency and oversight items ask whether teachers and students have adequate AI literacy, whether educators can choose whether to use GenAI, and whether students can opt out without disadvantage. Its diversity, non-discrimination and fairness items ask whether tools are accessible regardless of location, disability or socio-economic background, whether training-data bias is assessed, and whether institutions monitor negative impacts on particular groups.

From a policy perspective, the European Commission should support minimum expectations for equitable, accessible and agency-preserving GenAI deployment in higher education. These should include institutional access where GenAI is part of formal learning, accessibility-by-design, support for students with disabilities, guidance on premium tools, AI literacy support, bias monitoring, opt-out safeguards in high-stakes contexts, and a non-disadvantage principle for alternatives.

3.8 Data governance and procurement standards are not sufficiently adapted to educational LLM deployment

The deployment of LLMs and Generative AI in higher education creates data governance and procurement challenges only partly addressed by current EU rules. Data protection law remains fully applicable, and the AI Act adds requirements for AI systems and GPAI models (European Union, 2016; European Union, 2024). However, universities need operational guidance on procuring, approving, configuring and monitoring LLM-based systems in educational settings. The issue is sensitive because students and staff may enter personal data, coursework, examination answers, unpublished research, confidential institutional information, learning analytics, feedback, disciplinary records, accessibility-related information or other sensitive circumstances into GenAI systems. Student work may also be uploaded to external platforms for feedback, plagiarism analysis, assessment support, tutoring or administration, raising concerns about confidentiality, intellectual property, retention, reuse of prompts and outputs, cross-border transfers, vendor lock-in and model improvement.

The current framework sets general obligations but not sector-specific procurement standards for higher education. It does not define required vendor guarantees, pre-deployment information, excluded data categories, retention periods, reuse of inputs for training, update notification or monitoring of vendor terms. A mature approach should distinguish informal external use, approved institutional use and high-impact institutional deployment. Informal use concerns public tools used independently by students or staff; approved use concerns tools authorised for

defined purposes; high-impact deployment concerns systems integrated into teaching, assessment, student support, proctoring, admission, progression or administrative decision-making. Each level requires a different governance response.

For informal external use, institutions should specify which data must not be entered into public systems, including other students' personal data, confidential feedback, examination materials, unpublished research, sensitive institutional information, identifiable student work without authorisation and data subject to confidentiality or intellectual property restrictions. Students and staff should understand that public AI tools may process, retain or reuse inputs in opaque ways. For approved institutional use, universities should maintain an approved-tool list or equivalent mechanism, with review for data protection, cybersecurity, accessibility, contractual terms, pedagogical suitability and policy alignment. Approval should be tied to specific use cases and define permitted purposes, data categories, authorised users, student disclosure, support arrangements and incident reporting. High-impact deployments should undergo formal review covering data protection, cybersecurity, academic integrity, human oversight, appeal mechanisms, accessibility, bias, technical robustness, transparency, auditability and quality assurance integration.

Where the conditions established in Article 27 of the AI Act are met, the institution should conduct a Fundamental Rights Impact Assessment (FRIA) before deployment, according to the legal status of the institution and the nature of the deployment. This assessment should be coordinated with any required Data Protection Impact Assessment under Article 35 GDPR, while remaining analytically distinct: the FRIA addresses the wider effects on fundamental rights, whereas the DPIA focuses specifically on risks arising from the processing of personal data. The ADMIT legal taxonomy already identifies FRIA and DPIA as complementary instruments.

Procurement is therefore a central governance instrument. Vendor contracts should address data processing roles, retention and deletion, subcontractors, cross-border transfers, security, incident notification, continuity, model updates, access logs, audit rights, accessibility, user support, termination, data portability and restrictions on reuse of institutional or student data for training. Where GPAI models are used, vendors should identify the model provider, version, relevant documentation, known limitations, update policy and whether user data improve the model.

Data minimisation should be a core principle. Institutions should avoid sending personal or assessment-related data to LLM systems unless necessary for a defined purpose and supported by an appropriate legal and institutional basis. Where possible, data should be anonymised, pseudonymised or processed in controlled environments, but anonymisation should not be assumed where educational data can be re-identified through context, course details, writing style or metadata.

Informed consent should not be treated as a general or default legal basis for institutional GenAI processing. Universities must identify the appropriate Article 6 GDPR legal basis for each processing activity. Consent may be inappropriate where there is a power imbalance or where refusal or withdrawal would disadvantage the student; where consent is relied upon, it must be genuinely freely given, specific, informed and withdrawable.

Student work requires specific safeguards. Assignments, theses, code, portfolios and examination answers are academic artefacts that may contain personal expression, intellectual property, research ideas, sensitive information or evidence of learning. They should not be uploaded to external GenAI tools for feedback, evaluation or detection without institutional approval, appropriate safeguards and clear information to students.

Intellectual property and copyright also require policy attention. LLMs may generate teaching materials, adapt open educational resources, summarise copyrighted texts, produce images or code, and support research writing. Institutions need rules on output ownership, reuse, attribution, copyright compliance, third-party materials and infringement risks. The AI Act's GPAI copyright provisions are relevant, but universities still need practical guidance for educational content creation and reuse.

Cybersecurity and technical robustness should be built into procurement. LLM-based systems may be vulnerable to prompt injection, data leakage, unauthorised access, output manipulation, insecure integrations, external API dependency or service interruptions, especially when connected to learning management systems, student records, assessment platforms or administrative workflows. Procurement should therefore require evidence of security measures, vulnerability management, incident response, access control, logging and continuity arrangements.

The current EU framework lacks a higher education-specific procurement checklist. The AI Act and data protection law provide important foundations, but universities need guidance that combines them with academic integrity, quality assurance, student rights, accessibility, intellectual property and institutional governance. The European Commission should therefore support procurement and data governance guidance for LLMs and GenAI in higher education, including model contractual clauses, approved-tool procedures, data-handling rules, student-work safeguards, vendor documentation, update-notification duties, cybersecurity expectations, accessibility clauses and limits on reuse of student or institutional data. This would help institutions move from ad hoc adoption to responsible, auditable deployment.

3.9 GenAI governance is not yet embedded in European quality assurance mechanisms

The ESG 2015 provide the common European reference framework for quality assurance in higher education (ENQA et al., 2015). They cover internal and external quality assurance, student-centred learning, assessment, learning resources, student support, information management, public information, programme monitoring, periodic review and continuous improvement. These areas are directly affected by the adoption of LLMs and Generative AI systems. However, the ESG do not yet explicitly address GenAI governance, and this creates a significant implementation gap. The issue is not that the ESG are unsuitable for the GenAI context. On the contrary, they provide one of the most appropriate institutional channels for embedding responsible GenAI use into higher education practice. The ESG are designed to support both accountability and enhancement. They also recognise that quality in higher education depends on the interaction between teachers, students

and the institutional learning environment. LLMs now increasingly intervene in this interaction. They may support teaching design, student feedback, assessment preparation, tutoring, administrative communication, academic integrity procedures and student support. For this reason, they should become part of quality assurance, not remain outside it as a purely technical matter.

At present, many institutional responses to GenAI are likely to take the form of separate policy documents, guidance notes, acceptable-use rules or local experiments. These instruments are useful, but insufficient if they are disconnected from programme design, assessment policy, staff development, student support, public information and institutional monitoring. A university may publish a GenAI policy and still fail to review whether assessment tasks remain valid, whether students understand permitted uses, whether teachers have been trained, whether tools are accessible, whether AI-supported feedback is reliable, or whether students can appeal AI-mediated decisions.

Embedding GenAI governance into quality assurance means treating LLM adoption as part of the ordinary responsibility of higher education institutions for the quality of their provision. If GenAI tools affect learning opportunities, teaching methods, feedback processes, assessment evidence or student support, they must be reviewed through the same quality culture that governs curricula, learning outcomes, assessment methods and student experience. This approach would prevent GenAI from being governed only as an IT or compliance issue and would place it within the pedagogical mission of the institution.

Student-centred learning is a first area of direct relevance. GenAI may support personalisation, formative feedback, accessibility and self-paced learning. However, it may also encourage over-reliance, reduce independent reasoning, introduce opaque guidance or weaken the student's active role in learning. Quality assurance processes should therefore examine whether AI-enabled learning activities genuinely support student-centred learning, whether they preserve student agency, whether they are pedagogically justified, and whether they help students develop critical and independent capacities rather than simply obtain automated answers.

Assessment is a second critical area. LLMs challenge traditional assumptions about authorship, originality, process and evidence of learning. Quality assurance should therefore review whether assessment methods remain appropriate in the presence of GenAI, whether students are clearly informed about permitted and prohibited uses, whether disclosure requirements are consistent, and whether assessment tasks are aligned with intended learning outcomes. AI-supported assessment practices should also be reviewed for reliability, fairness, human oversight, contestability and transparency.

Programme design and review are also affected. If GenAI becomes part of teaching and learning, programme teams should consider where and how AI literacy is developed, how disciplinary uses of GenAI are addressed, and how students are prepared to use AI critically and responsibly in academic and professional contexts. Programme review should also consider whether GenAI changes the skills that courses aim to develop, the forms of evidence used to assess learning, and the support needed by students and staff. Learning resources and student support represent

another area where GenAI requires quality assurance attention. AI tutors, chatbots, writing assistants, administrative helpdesks and wellbeing-oriented tools may extend support capacity and provide timely assistance. However, they may also give inaccurate advice, fail to recognise complex student needs, process sensitive data, or create confusion about when human support is available. Quality assurance should therefore examine whether AI-supported services are reliable, accessible, transparent, appropriately supervised and connected to human escalation routes.

The table below illustrates how GenAI governance could be mapped onto existing ESG-related quality assurance areas.

Table 3. Mapping GenAI Governance to ESG-Oriented Quality Assurance Areas

ESG-oriented quality assurance area	GenAI-related governance question	Possible institutional evidence
Policy for quality assurance	Is GenAI governance included in the institution's quality policy, academic integrity policy and digital transformation strategy?	Institutional GenAI policy, governance framework, approval procedures, responsibility matrix
Programme design and approval	Are GenAI competences, AI literacy and disciplinary implications considered in programme design?	Programme specifications, curriculum mapping, AI literacy learning outcomes, programme approval records
Student-centred learning and teaching	Does GenAI support active learning, student agency and pedagogically justified innovation?	Course designs, teacher guidance, student feedback, evaluation of AI-supported activities
Assessment of students	Are assessment methods valid, fair and transparent in the presence of GenAI?	Assessment policies, disclosure rules, redesigned assignments, human review procedures, appeal records
Teaching staff	Are educators trained to use GenAI responsibly and to supervise AI-supported learning and assessment?	CPD records, teacher training materials, guidance documents, staff support mechanisms
Learning resources and student support	Are AI-supported services reliable, accessible, supervised and connected to human support?	Approved tool lists, accessibility reviews, escalation procedures, service monitoring reports
Information management and review	Does the institution collect evidence on GenAI use, risks, incidents and improvement needs?	Self-assessment results, incident logs, student surveys, complaints analysis, risk reviews, updated policies

Source: authors' elaboration based on ESG 2015 and the ADMIT Institutional Self-Assessment Tool (ENQA et al., 2015; Triantafyllopoulos et al., 2025a).

This mapping shows that GenAI governance does not require a separate quality assurance framework, but the systematic interpretation of existing ESG areas in light of LLM-related risks and opportunities. The ESG already provide the institutional logic of implementation, monitoring, review

and improvement; what is missing is explicit EU-level guidance on integrating GenAI into these processes.

The ADMIT Institutional Self-Assessment Tool can support this integration by translating the ethical taxonomy into role-specific questions for students, educators and institutional leadership. Its results can feed into internal quality assurance, programme review, staff development, student support planning and policy updates, linking ethical reflection with evidence-based governance.

A quality assurance approach can also accommodate Europe's diverse institutional contexts. Because the ESG are deliberately generic, they allow common principles to be adapted to different systems, missions, delivery modes and resources without losing coherence. External quality assurance may also play a role, provided it avoids rigid checklists and instead reviews whether institutions have appropriate policies, safeguards, monitoring and improvement processes for GenAI in learning and teaching. The gap is therefore clear: the ESG provide the quality assurance architecture, ADMIT offers ethical and operational tools, and the AI Act provides the legal baseline. What remains missing is a sector-specific EU policy layer connecting these instruments to programme review, assessment design, student-centred learning and institutional improvement.

From a policy perspective, the European Commission should encourage the integration of GenAI governance into ESG-aligned internal and external quality assurance, through guidance developed with quality assurance agencies, higher education institutions, student organisations and European networks. The aim should be to make GenAI governance part of a continuous cycle of institutional responsibility, monitoring, stakeholder involvement and improvement.

3.10 Lack of systematic monitoring and iterative policy updating

The current EU framework lacks a dedicated mechanism for monitoring the evolution of LLM and Generative AI use in higher education. This is a significant gap because technologies, markets, institutional practices and student behaviour are changing rapidly through new model releases, multimodal functions, platform integrations, pricing models and patterns of use. Static guidance is therefore unlikely to remain adequate. Monitoring should go beyond legal compliance with the AI Act, data protection and procurement rules, and also capture pedagogical effectiveness, academic integrity, student experience, teacher workload, accessibility, bias, equity, stress, over-reliance, privacy incidents, appeals, technical failures, environmental impact and the quality of AI-supported learning.

Many GenAI risks become visible only after deployment. Tools that appear useful in procurement or pilots may reveal limitations across disciplines, languages, student groups, assessment formats or institutional contexts. Bias, accessibility problems, academic integrity issues and staff workload may emerge only during regular use. Monitoring must therefore be longitudinal, covering adoption, updates, scaling, repurposing and connection to high-impact processes. It should draw on evidence from students, educators, administrative staff, quality assurance units, data protection officers, IT services and institutional leadership, as well as incidents, complaints, appeals, accessibility reports, integrity cases, training needs and pedagogical value.

The ADMIT Institutional Self-Assessment Tool offers a practical model for such monitoring (Triantafyllopoulos et al., 2025a). By translating the ADMIT taxonomy into role-specific questionnaires for students, educators and institutional leadership, it captures different perspectives on the same system: strategy, procurement and risk management for leaders; assessment, workload and pedagogical value for teachers; and clarity, fairness, access, trust and opt-out for students.

A monitoring framework could build on this multi-stakeholder logic through an annual or periodic GenAI self-assessment cycle structured around the eight ADMIT dimensions: educational impact and academic integrity; privacy and data governance; wellbeing; agency and oversight; diversity, non-discrimination and fairness; accountability; transparency; and technical robustness and safety. Results should inform action plans, quality assurance review, staff development, procurement and policy updates. Monitoring should also be proportionate to risk: optional tools may require light-touch review, institutionally provided tools structured monitoring, and systems used in assessment, admission, proctoring, progression, integrity investigations or profiling more formal and frequent review, including fairness and contestability. Institutional monitoring alone is insufficient because GenAI systems operate across borders, platforms and providers, and similar risks may appear across universities. Fragmented monitoring may miss recurring European-level problems, such as false positives in proctoring, bias in language feedback, privacy concerns or accessibility barriers. EU policy should therefore support sectoral evidence gathering through existing bodies and networks, including the AI Office, higher education networks, quality assurance agencies, digital education platforms and sectoral observatories. The aim should be to identify emerging risks, document good practices, support policy learning and update guidance promptly.

Iterative policy updating is essential. EU guidance on GenAI in higher education should not be a one-off document, but should be reviewed periodically in light of technological developments, institutional experience, legal interpretation, stakeholder feedback and evidence of harms or benefits. Updates should respond to documented incidents, emerging inequalities, new assessment practices, provider behaviour, new student uses and evidence from self-assessment and quality assurance.

From a policy perspective, the European Commission should encourage a periodic GenAI monitoring cycle for higher education institutions, aligned with the ADMIT taxonomy and integrated into quality assurance. It should also support European-level evidence aggregation, exchange of practice and regular guidance updates. This would keep GenAI governance adaptive, evidence-based and responsive to institutional experience, rather than limited to static compliance at adoption.

3.11 Synthesis: the need for a sector-specific EU policy layer

The gaps identified in this subsection point to a common conclusion. The current European framework already contains several essential components for governing LLMs and Generative AI in higher education, but these components are not yet connected through a sector-specific policy layer. The EU AI Act provides the legal baseline for risk classification, GPAI obligations, high-risk systems and transparency requirements. The Digital Education Action Plan provides the strategic framework for inclusive, accessible and high-quality digital education. The ESG 2015 provide the quality

assurance architecture for learning, teaching, assessment, student support and continuous improvement. The ADMIT outputs provide the ethical and operational bridge through the taxonomy, the guide and the institutional self-assessment tool. The problem is therefore not the absence of relevant European instruments. The problem is that each instrument addresses only part of the governance challenge. The AI Act is horizontal and risk-based; it does not define the internal academic procedures through which universities should govern GenAI. The Digital Education Action Plan supports digital transformation and competences; it does not yet provide detailed operational guidance for the current generation of LLMs. The ESG establish quality assurance principles; they do not yet explicitly translate those principles into GenAI-related institutional requirements. The ADMIT outputs provide a strong ethical and practical basis; they need to be connected to EU-level policy guidance and institutional implementation mechanisms.

This creates a need for a policy layer that is neither a completely new regulatory regime nor a purely voluntary set of generic recommendations. Such a layer should translate existing EU principles into concrete expectations for higher education institutions. It should clarify how universities classify LLM use cases, how they distinguish low-impact assistive uses from high-impact academic decision-making uses, how they disclose AI deployment to students and staff, how they assign responsibility, how they protect academic integrity, how they ensure human oversight and appeal, how they govern student data, how they guarantee equitable access, and how they embed GenAI into quality assurance.

A sector-specific policy layer would also reduce fragmentation across the European Higher Education Area. Without common guidance, institutions may adopt very different approaches to GenAI: some may prohibit it broadly, others may integrate it rapidly, others may leave decisions to individual teachers or departments, and others may rely on vendor-led solutions. Such diversity can support innovation, but it can also create uncertainty, unequal student protection, inconsistent academic integrity rules and uneven institutional readiness. EU-level guidance would not eliminate institutional autonomy; it would provide a shared baseline that institutions can adapt to their own legal, pedagogical and organisational contexts.

The synthesis of the gaps also shows that GenAI governance must be multi-dimensional. Transparency without accountability is insufficient, because affected students must know not only that AI is used, but also who is responsible and how to challenge outcomes. Accountability without technical robustness is weak, because institutions cannot responsibly deploy systems that are unreliable, insecure or unstable. Technical robustness without academic integrity is incomplete, because a technically reliable system may still undermine the validity of assessment if used inappropriately. Equity without opt-out safeguards may still leave students exposed to mandatory AI-mediated processes. Quality assurance without monitoring may fail to detect emerging harms or ineffective practices.

The ADMIT ethical taxonomy is particularly valuable because it captures this multi-dimensional character. Its eight dimensions make clear that responsible GenAI governance in higher education must address educational impact and academic integrity, privacy and data governance, societal, individual and environmental wellbeing, teacher and student agency and oversight,

diversity, non-discrimination and fairness, accountability, transparency, and technical robustness and safety. These dimensions provide a structured way to move from isolated concerns to an integrated institutional framework.

At the same time, the taxonomy should not remain only a conceptual reference. The D5.3 Guide translates the ethical dimensions into practical guidance for educators, students and institutional leaders. The D5.4 Institutional Self-Assessment Tool further operationalises them through role-specific reflection and feedback. Together, these outputs demonstrate that ethical GenAI governance can be made actionable. They also show that institutions can assess their own maturity, identify gaps, and connect ethical reflection to policy design and quality assurance.

The policy task is therefore to connect legal compliance, ethical governance and quality assurance into a single implementation logic. A university adopting GenAI should be able to answer a coherent set of questions: What is the intended purpose of the system? Does the use case fall within a low-impact, transparency-sensitive, high-impact or potentially high-risk area? What information has been obtained from the provider? What data are processed? Who is responsible for procurement, deployment and monitoring? How are students and teachers informed? What human oversight is provided? Can affected students appeal? Are the tools accessible and equitably available? How is the practice reviewed through quality assurance?

The following cross-reference matrix summarises how the main European and ADMIT references can be jointly mobilised to support such a policy layer. Its purpose is to make explicit the relationship between regulatory provisions, ethical dimensions and institutional implications for higher education.

Table 4. Cross-reference between EU regulatory provisions, ADMIT ethical dimensions and policy implications for higher education

EU / ADMIT reference	Main issue for higher education	Relevant ADMIT ethical dimension(s)	Policy implication
EU AI Act, risk-based approach	The regulatory classification of LLM-based systems in higher education depends on the concrete use case, not on the mere use of a GenAI model. Generic study support, teaching assistance, assessment, grading, admission, proctoring and academic progression raise different levels of risk.	Accountability; Transparency; Technical Robustness and Safety	EU guidance should clarify how higher education institutions should classify different educational uses of LLMs, distinguishing low-impact assistive uses from AI-mediated academic decisions.
EU AI Act, Articles 51–55 on GPAI models	GPAI obligations mainly concern model providers, while universities often act as deployers integrating	Transparency; Accountability	EU policy should translate GPAI provider obligations into operational duties for

	third-party LLMs into teaching, learning, assessment or student-support processes.		higher education institutions, especially concerning disclosure, documentation, procurement, monitoring and institutional responsibility.
EU AI Act, Annex III on high-risk AI systems	Some educational uses of AI may fall within the high-risk category when they affect access to education, admission, evaluation, allocation, progression or other significant academic outcomes.	Accountability; Technical Robustness and Safety; Diversity, Non-discrimination and Fairness	AI-mediated academic decisions should require human oversight, contestability, auditability, impact assessment, documented reasoning and accessible appeal procedures.
ADMIT ethical taxonomy	The ADMIT taxonomy identifies key dimensions for the ethical use of GenAI in higher education, including educational impact and integrity, privacy, wellbeing, agency, fairness, accountability, transparency and technical robustness.	All ADMIT dimensions	EU policy should adopt a sector-specific ethical framework for GenAI in higher education, ensuring that legal compliance is complemented by pedagogical, institutional and student-centred safeguards.
ADMIT D5.3 Guide of ethical use of LLM / GenAI systems in higher education	Responsible GenAI adoption requires role-specific guidance for educators, students and institutional leadership, including disclosure, accuracy checking, human oversight, privacy protection, equitable access and academic integrity safeguards.	Educational Impact and Academic Integrity; Teacher and Student Agency and Oversight; Transparency; Accountability; Diversity, Non-discrimination and Fairness	EU recommendations should define minimum institutional standards for GenAI use in teaching, assessment, student support and academic governance.
ADMIT D5.4 Institutional Self-Assessment Tool	The ADMIT taxonomy has been operationalised into role-specific self-assessment processes for students, educators and institutional leadership. This creates a bridge between ethical principles	Accountability; Transparency; Technical Robustness and Safety; all ADMIT dimensions	Higher education institutions should be encouraged to conduct periodic GenAI self-assessment and use the results for quality assurance, risk management and

Digital Education Action Plan 2021–2027	and institutional monitoring.		continuous improvement.
	The Action Plan promotes high-quality, inclusive and accessible digital education, digital skills and AI-related competences, but it does not provide a detailed operational framework for the current generation of LLM-based educational systems.	Teacher and Student Agency and Oversight; Diversity, Non-discrimination and Fairness; Societal, Individual and Environmental Wellbeing	The Action Plan should be complemented by updated EU guidance specifically addressing GenAI and LLMs in higher education, with emphasis on AI literacy, inclusion, transparency and institutional readiness.
ESG 2015	The ESG provide the European reference framework for quality assurance in higher education, including student-centred learning, assessment, programme review, public information, stakeholder involvement and continuous improvement.	Accountability; Educational Impact and Academic Integrity; Transparency	GenAI governance should be embedded into existing internal and external quality assurance processes, rather than treated as a separate or purely technical policy area.
Combined regulatory and ethical gap	The AI Act establishes general legal obligations, while ADMIT provides an ethical and operational perspective for higher education. However, a sector-specific EU policy layer is still needed to connect regulation, institutional governance and academic practice.	Accountability; Transparency; Technical Robustness and Safety; Educational Impact and Academic Integrity	The European Commission should issue sector-specific recommendations for the ethical and responsible use of LLMs in higher education, aligned with the AI Act, ADMIT outputs, the Digital Education Action Plan and ESG 2015.

Source: authors' elaboration based on Regulation (EU) 2024/1689, the Digital Education Action Plan 2021–2027, ESG 2015, and the ADMIT outputs D5.2, D5.3 and D5.4 (European Union, 2024; European Commission, 2020; ENQA et al., 2015; Carbonel et al., 2024; Triantafyllopoulos et al., 2025a; Triantafyllopoulos et al., 2025b).

The cross-reference matrix shows that the governance of LLMs and Generative AI in higher education cannot be reduced to technical compliance with the AI Act. Legal obligations concerning GPAI models, high-risk systems and transparency need to be connected with academic integrity, student rights, teacher responsibilities, institutional accountability, equitable access, technical robustness and quality assurance. This connection is precisely where EU policy can add value.

The recommendations that follow are therefore designed as a sector-specific translation of existing European frameworks. They do not assume that higher education requires a completely

separate AI regulatory regime. Rather, they propose that the European Commission provide practical guidance, minimum standards and monitoring mechanisms capable of helping institutions implement the AI Act and related EU policy instruments in ways that are appropriate to the academic context.

4 Concrete Recommendations to the European Commission

The analysis developed in the previous subsections shows that the ethical and responsible use of LLMs and Generative AI systems in higher education cannot be governed only through general AI regulation or broad digital education strategies. The AI Act provides the legal baseline, but higher education requires an additional policy layer capable of translating legal obligations into academic procedures, institutional responsibilities, student-facing safeguards and quality assurance mechanisms.

The recommendations below are addressed to the European Commission and are intended to support a coherent European approach to GenAI governance in higher education. They do not propose the creation of a separate regulatory regime for universities. Rather, they aim to make existing European instruments operational in the academic context, aligning the AI Act, the Digital Education Action Plan, ESG 2015 and the ADMIT outputs (European Union, 2024; European Commission, 2020; ENQA et al., 2015; Carbonel et al., 2024; Triantafyllopoulos et al., 2025a; Triantafyllopoulos et al., 2025b).

In accordance with Article 165 TFEU, which preserves Member State responsibility for the organisation and content of education systems, the proposed European policy layer should primarily take the form of guidance, Council recommendations, coordination mechanisms, funding incentives, common tools and support for mutual learning, rather than harmonisation of national higher education law: EU action in education supports, coordinates and supplements Member State action. Moreover, since the European Higher Education Area extends beyond the European Union, recommendations intended to apply across the EHEA as a whole should be implemented not only through European Commission instruments but also through the Bologna Follow-Up Group, quality assurance networks and the relevant EHEA stakeholder organisations.

Methodological basis. *The recommendations below were derived through a structured analytical process. First, the eight ADMIT ethical dimensions (D5.2; Carbonel et al., 2024) were mapped onto the risk-based classification logic of the EU AI Act, identifying where existing legal provisions require sector-specific interpretation for higher education (§2). Second, a systematic gap analysis (§3) identified ten areas where EU regulation and policy do not yet provide sufficient operational guidance for universities. Third, these gaps were cross-referenced with the practical guidance developed in the ADMIT Guide of ethical use (D5.3; Triantafyllopoulos et al., 2025b) and with the role-specific questions of the ADMIT Institutional Self-Assessment Tool (D5.4; Triantafyllopoulos et al., 2025a), ensuring that each recommendation is grounded in both the European regulatory landscape and the empirical and ethical work of the ADMIT partnership. The*

recommendations were further informed by the institutional practices mapped within the ADMIT consortium (WP3), the focus groups and pilot trials conducted within WP4, and the peer learning activities carried out within WP6. Each recommendation below specifies: the regulatory gap it addresses, the ADMIT ethical dimensions it operationalises, the stakeholder groups primarily concerned and the type of EU-level action proposed.

Stakeholder framework. *As mandated in the ADMIT Description of the Action, the recommendations are structured around the three key stakeholder perspectives in higher education: (a) students, who need clear rules, disclosure, appeal mechanisms, equitable access and AI literacy; (b) teachers, who need operational guidance, training, assessment redesign support and meaningful oversight responsibilities; and (c) institutions (including leadership, administrative and support layers), which need governance frameworks, procurement standards, classification procedures, quality assurance integration and monitoring mechanisms. Each recommendation specifies stakeholder implications and ADMIT ethical dimensions, showing how it extends or contextualises the AI Act for the higher education sector.*

4.1 Establish EU-level guidance for the classification and deployment of LLM-based systems in higher education

The European Commission should develop sector-specific guidance to help higher education institutions classify LLM-based systems according to their intended purpose, institutional context and impact on students, educators and applicants (European Union, 2024; Carbonel et al., 2024). The present document, together with the ADMIT ethical taxonomy (D5.2), the Guide of ethical use of LLM / Generative AI systems in Higher Education (D5.3) and the Institutional Self-Assessment Tool (D5.4), is intended to serve as a substantive basis for such guidance, offering classification criteria, ethical dimensions and operational tools developed through a European cooperation partnership of distance and open universities. This guidance should clarify that the legal and ethical treatment of an LLM depends not on the model alone, but on how it is used within academic processes.

The guidance should distinguish between optional individual study support, teacher-mediated pedagogical use, institutionally provided assistive tools, transparency-sensitive student interaction systems, and high-impact academic decision-support systems. It should also clarify when educational uses may fall within the high-risk logic of Annex III of the AI Act, particularly in relation to admission, assessment, grading, allocation of opportunities, academic progression, proctoring and detection of prohibited behaviour during tests.

Such guidance should help institutions avoid both under-classification and over-classification. Under-classification occurs when high-impact systems are treated as ordinary digital tools simply because they rely on a general-purpose model or because a human remains formally present. Over-classification occurs when all LLM uses are treated as high-risk, including optional and low-impact pedagogical uses that can be governed through lighter safeguards. A balanced classification framework would protect students while preserving space for responsible pedagogical innovation.

The Commission should also provide templates for institutional classification records. These records should document the intended purpose of the system, the users involved, the affected persons, the data processed, the role of AI outputs in the workflow, the possible impact on academic outcomes, the human oversight arrangements, and the available appeal or review mechanisms.

ADMIT dimensions: Accountability, Transparency, Technical Robustness and Safety. AI Act extension: operationalises the risk-based classification logic of Article 6 and Annex III, together with the relevant transparency obligations under Article 50, for higher education use cases, where the AI Act does not specify how universities should classify educational GenAI deployments. For students: clear information on how each AI tool used in their educational experience has been classified and what safeguards apply. For teachers: practical guidance on which tools require heightened scrutiny and which may be used with lighter governance. For institutions: a documented classification process covering all GenAI deployments, with templates for classification records.

4.2 Define a common transparency standard for educational GenAI use

The European Commission should support the development of a common transparency standard for the use of LLMs and Generative AI systems in higher education, in cooperation with higher education institutions, quality assurance agencies and stakeholder networks (European Union, 2024; Carbonel et al., 2024). While the EU does not directly regulate higher education governance, it can provide guidance and reference frameworks that help institutions adopt coherent transparency practices. Such a standard should move beyond generic disclosure and specify what students, educators and institutional stakeholders need to know in different academic contexts.

At minimum, students should be informed when an AI system is used in teaching, learning support, assessment, feedback, proctoring, academic integrity procedures, student profiling, admission or administrative guidance. The information should explain the purpose of the system, the nature of the AI-generated output, whether the output is advisory or decision-shaping, whether personal or assessment-related data are processed, what limitations are known, and how human review or redress can be requested.

Educators should receive more detailed operational information. They should know which systems are institutionally approved, what uses are permitted, what data may be entered, what outputs require verification, how AI-supported teaching or assessment should be disclosed to students, and how to report problems. Institutional leaders and quality assurance bodies should receive documentation suitable for procurement, monitoring, auditability and policy review.

The Commission should encourage institutions to maintain a register of approved GenAI systems used in formal academic contexts. Such a register should not necessarily disclose confidential technical or security information, but it should provide sufficient information for students, staff and governance bodies to understand where AI is used and under what conditions.

Transparency should be proportionate to risk. Optional study-support tools may require lighter notices and responsible-use guidance. Systems used in assessment, admission, proctoring, student profiling or academic integrity procedures require more detailed disclosure, documentation

and procedural safeguards. The key principle should be that affected persons must be able to understand the role of AI in the academic process and the safeguards available to them.

ADMIT dimensions: Transparency, Accountability. AI Act extension: specifies the transparency and disclosure requirements of the AI Act for educational contexts, where generic ‘AI interaction’ notices are insufficient for students who need to understand the role of AI in their assessment, feedback, support or administrative processes. For students: disclosure notices whenever AI is used in teaching, assessment, support or administrative processes affecting them; accessible explanations of AI role and limitations. For teachers: information on the capabilities, limitations and known biases of AI tools they are expected to use. For institutions: institutional disclosure policies, model AI-use notices, and procedures for communicating AI deployment to all affected parties.

4.3 Establish minimum safeguards for AI-mediated academic decisions

The European Commission should support the definition of minimum safeguards for AI-mediated academic decisions (European Union, 2024; Triantafyllopoulos et al., 2025b). These safeguards should apply whenever an LLM-based or GenAI-enabled system contributes to assessment, grading, admission, allocation of academic opportunities, academic progression, proctoring, misconduct detection, student-risk profiling or other decisions that may significantly affect students or applicants.

The first safeguard should be meaningful human oversight. Significant academic decisions should not be based solely on AI outputs. Human reviewers should have the competence, time, authority and information required to understand, challenge, correct or disregard AI-generated recommendations, scores, flags or summaries. Oversight should not be reduced to formal approval of automated outputs.

The second safeguard should be contestability. Students and applicants should have access to clear and timely procedures for challenging AI-mediated outcomes. Where an AI system has contributed to a decision or flag, the affected person should be informed of this contribution, receive a meaningful explanation of the role played by the system, have access to relevant evidence where appropriate, and be able to request human review by a competent decision-maker.

The third safeguard should be evidentiary caution. AI-generated suspicion, risk scores, proctoring alerts or misconduct flags should never be treated as conclusive evidence. They should be considered only as signals requiring contextual human review. This is particularly important in academic integrity and proctoring contexts, where false positives may cause reputational harm, anxiety and procedural unfairness.

The fourth safeguard should be documentation. Institutions should document the role of AI systems in academic decisions, the review process followed, the reasons for the final decision, and any appeal or correction measures. This documentation should support accountability, quality assurance and continuous improvement.

ADMIT dimensions: Teacher and Student Agency and Oversight, Accountability, Educational Impact and Academic Integrity. AI Act extension: translates Article 14 human oversight requirements into the specific context of academic decision-making, where AI outputs may influence grades, misconduct findings, progression or admission without adequate human review. For students: accessible appeal and contestability mechanisms for AI-mediated or AI-influenced academic decisions; information on how to challenge AI-generated flags or scores. For teachers: meaningful review authority with sufficient time, competence and information to assess, override or reject AI outputs. For institutions: documented human oversight procedures for all high-impact GenAI deployments; formal appeal routes; periodic audit of actual human review practices.

4.4 Develop common European reference principles for institutional academic integrity policies on GenAI

The European Commission should encourage the development of common European reference principles for institutional academic integrity policies concerning LLMs and Generative AI (Triantafyllopoulos et al., 2025b). These standards should recognise that academic integrity is not only a misconduct issue, but a core condition for the validity of learning, assessment, certification and trust in higher education.

The standards should require institutions to define permitted, restricted and prohibited uses of GenAI in coursework, examinations, research-related activities, coding tasks, language support, collaborative work and student support. These rules should be understandable to students, adaptable to disciplinary contexts and communicated clearly at course and programme level.

Disclosure should be treated as a central academic integrity mechanism. Students should know when GenAI use must be declared, how it should be declared, and how disclosure will be evaluated. Disclosure should not be framed only as a punitive tool. It can also support reflection, transparency and responsible learning, especially where students document how they used AI, how they verified outputs and how they transformed AI-generated material into their own work.

The Commission should also discourage over-reliance on AI-detection systems. Detection tools should not be used as automatic evidence of misconduct, and students should not be sanctioned solely on the basis of probabilistic or opaque AI-generated indicators. Institutions should establish fair procedures including human review, access to relevant evidence, the possibility for students to respond, and proportional decision-making.

Finally, academic integrity standards should encourage assessment redesign. Institutions should support educators in developing assessment formats that remain valid in the presence of GenAI, including oral components, in-person or online supervised examination, process portfolios, supervised stages of work, reflective commentaries, authentic tasks, iterative drafts and critical evaluation of AI outputs. This would shift the policy response from surveillance to better assessment design. Common European reference principles for assessment redesign should be flexible enough to accommodate the full range of European higher education delivery models. Distance higher education institutions, whose students are geographically dispersed and whose assessment is

conducted online or through a combination of remote and periodic in-person sessions, have developed valid formats - including AI-proctored remote examinations, synchronous oral assessments via videoconference, and supervised in-person sessions where required - that should be explicitly recognised as equivalent alternatives. Guidance that implicitly treats in-person supervised examination as the only credible integrity safeguard would be both technically incorrect and institutionally *discriminatory*.

ADMIT dimensions: Educational Impact and Academic Integrity, Transparency. AI Act extension: fills a gap where the AI Act does not address academic integrity as a distinct regulatory category, yet GenAI profoundly affects assessment validity, authorship norms, misconduct detection and scholarly honesty in higher education. For students: clear rules on permitted, restricted and prohibited uses of GenAI in coursework, examinations and assessments; fair misconduct procedures with human review; protection from false AI-detection accusations. For teachers: support for assessment redesign in GenAI-rich environments; guidance on integrating AI tools constructively while preserving authentic learning. For institutions: comprehensive academic integrity policies with disclosure norms, assessment redesign, GenAI-specific misconduct procedures and periodic review.

4.5 Ensure equitable access, accessibility and opt-out safeguards

The European Commission should promote minimum expectations for equitable and accessible GenAI deployment in higher education (Carbonel et al., 2024; Triantafyllopoulos et al., 2025a). Where GenAI tools are formally integrated into teaching, learning or assessment, institutions should ensure that students are not disadvantaged by unequal access to commercial tools, premium subscriptions, devices, connectivity, language support or AI literacy.

Formal learning activities should not assume that all students have equivalent private access to advanced LLMs. If the use of GenAI is expected or required, institutions should provide access under fair and transparent conditions or offer equivalent alternatives. This is particularly important where GenAI use may influence assessment performance, feedback quality, participation in learning activities or access to support.

Accessibility should be treated as a design requirement. GenAI systems used in formal educational processes should be assessed for compatibility with assistive technologies, usability for students with disabilities, availability of alternative interaction modes, clarity of interface, multilingual needs where relevant, and access to human support. Institutions should also monitor whether GenAI systems produce unequal effects for specific groups of students.

Accessibility reviews should be based on applicable European legislation and recognised technical standards, including Directive (EU) 2016/2102, Directive (EU) 2019/882 and EN 301 549, where relevant. Web-based GenAI interfaces should also be evaluated against the applicable WCAG success criteria, using both technical testing and user evaluation involving persons with disabilities. Procurement and approval procedures should require demonstrable and testable conformity rather than treating accessibility-by-design as an abstract principle.

The Commission should also encourage opt-out safeguards in high-stakes contexts. Students should have access to alternative procedures where GenAI use may affect assessment, proctoring, academic integrity investigations, student profiling, privacy, accessibility or dignity. Opt-out should not become a generic veto over all digital tools, but it should be available where AI-mediated processes may significantly affect student rights or academic opportunities.

The non-disadvantage principle should guide these alternatives. A student who uses a legitimate opt-out procedure should not be penalised with a less favourable assessment format, excessive administrative burden, delayed progression or reduced access to support, unless a difference is objectively justified and proportionate.

ADMIT dimensions: Diversity, Non-discrimination and Fairness, Societal, Individual and Environmental Wellbeing. AI Act extension: specifies the non-discrimination, accessibility and fundamental rights provisions of the AI Act for the education sector, ensuring equitable access to GenAI tools regardless of disability, location, language, socio-economic background or institutional delivery model. For students: equitable access to GenAI tools in formal learning activities; opt-out mechanisms that do not disadvantage non-users; accessible alternatives for students with disabilities. For teachers: awareness of bias risks in AI outputs and assessment; training on equitable deployment. For institutions: accessibility-by-design procurement standards; monitoring of unequal effects on specific student groups; equitable access policies covering distance and on-campus learners.

4.6 Ensure environmentally proportionate GenAI deployment

The European Commission should encourage higher education institutions to assess the environmental proportionality of institutional GenAI use. Procurement and deployment decisions should consider whether the educational benefit justifies the computational and material resources required, whether smaller or locally controlled models could fulfil the same purpose, and whether suppliers provide sufficiently transparent information on resource consumption. Environmental indicators should be incorporated into procurement, periodic self-assessment and quality assurance, particularly for mandatory, large-scale or continuously operating services. This recommendation operationalises the ADMIT dimension of Societal, Individual and Environmental Wellbeing for the higher education sector.

4.7 Make structured AI literacy and teacher training a precondition for formal institutional deployment

The European Commission should make structured AI literacy and teacher training a core element of responsible GenAI deployment in higher education. The AI Act recognises the importance of AI literacy, and the higher education sector requires this principle to be operationalised for students, educators, institutional leaders and administrative staff (European Union, 2024).

Teacher training is particularly important, especially if AI literacy is understood not as simple technical training, but as critical awareness of system functioning, risks, output interpretation and

responsible use, as also emphasised in recent national implementation initiatives on AI (Presidency of the Council of Ministers, 2026). Educators are expected to make pedagogical decisions, verify AI-generated outputs, redesign assessment, guide students, protect academic integrity, and participate in human oversight. These tasks require more than basic familiarity with digital tools. Teachers need to understand how LLMs generate outputs, why hallucinations occur, how bias may appear, what data should not be entered into external systems, how to disclose AI-supported teaching practices, and how to evaluate student work in GenAI-rich environments.

Training should also address automation bias, over-reliance, false positives in AI detection, limitations of proctoring systems, accessibility, data protection, intellectual property, prompt design, verification strategies, and the ethical implications of AI-mediated feedback and assessment. For staff involved in admission, progression, misconduct procedures or student support, training should include procedural fairness, evidence evaluation, student rights and appeal mechanisms.

Students also require structured AI literacy. They need to understand permitted and prohibited uses, disclosure expectations, risks of hallucination, responsible prompting, verification of outputs, data protection, academic integrity, bias and the limits of AI-generated advice. AI literacy should not be treated only as a technical competence, but as part of critical academic practice.

The Commission should encourage institutions to connect AI literacy with curriculum design, staff development, academic integrity policies and quality assurance. Formal institutional deployment of GenAI tools should be accompanied by training and support for the users affected by the deployment.

ADMIT dimensions: Teacher and Student Agency and Oversight, Educational Impact and Academic Integrity. AI Act extension: operationalises the AI Act's AI literacy requirements (Article 4) for higher education, where both students and teachers need structured competence development to use, evaluate and critically engage with GenAI systems. For students: structured AI literacy programmes as part of academic practice, covering critical evaluation, verification, responsible use and ethical awareness. For teachers: professional development connecting AI literacy with pedagogical practice; training on output verification, hallucination risks, bias, assessment redesign and human oversight responsibilities. For institutions: AI literacy and teacher training as institutional prerequisites before formal GenAI deployment; resource allocation for ongoing professional development.

4.8 Introduce periodic institutional self-assessment and EU-level monitoring

The European Commission should encourage higher education institutions to conduct periodic self-assessment of GenAI use, aligned with ethical, legal and quality assurance principles (Triantafyllopoulos et al., 2025a; ENQA et al., 2015). The ADMIT Institutional Self-Assessment Tool provides a strong operational basis for this recommendation, because it translates ethical dimensions into role-specific reflection for students, educators and institutional leadership.

Periodic self-assessment should help institutions identify how GenAI is being used, what risks emerge, whether policies are understood, whether students and staff feel adequately supported, whether tools are accessible, whether complaints or appeals are arising, whether academic integrity rules are effective, and whether AI-supported practices are improving or weakening the quality of learning.

Self-assessment should not remain a symbolic exercise. Results should feed into institutional action plans, procurement decisions, staff development, student support, assessment redesign, risk management and quality assurance review. Institutions should be encouraged to document improvement measures and review them periodically.

For high-impact deployments, self-assessment should be complemented by outcome-based monitoring. Relevant indicators may include error and false-positive rates, frequency and outcomes of appeals, evidence of automation bias, disaggregated effects across student groups, accessibility incidents, effects on learning outcomes, staff workload, data-protection incidents and environmental resource use. Institutions should establish baselines, review thresholds and escalation criteria, with independent or external evaluation where the potential impact on rights or academic outcomes is substantial.

At EU level, the Commission should support the aggregation of evidence on GenAI deployment in higher education. This could be done through existing networks, quality assurance agencies, European higher education bodies, digital education platforms and sectoral observatories. The purpose should be to identify emerging risks, document good practices, monitor inequalities, analyse incidents, and update guidance in light of technological and institutional developments.

GenAI governance should be iterative. Guidance should be reviewed periodically in response to changes in model capabilities, provider practices, assessment methods, student behaviour, institutional experience and evidence of harm or benefit. A static policy framework would quickly become outdated in a rapidly evolving technological environment.

ADMIT dimensions: Accountability, all eight dimensions via the ADMIT Institutional Self-Assessment Tool (D5.4). AI Act extension: creates a sector-specific monitoring and self-assessment framework that goes beyond the AI Act's requirements, addressing the need for ongoing institutional learning, iterative policy updating and EU-level evidence gathering on GenAI in higher education. For students: feedback channels and involvement in institutional GenAI review processes. For teachers: participation in self-assessment, reporting of incidents and contribution to institutional learning. For institutions: periodic self-assessment cycles structured around the eight ADMIT dimensions; documented improvement plans; EU-level monitoring through existing higher education networks and quality assurance agencies.

4.9 Embed GenAI governance into ESG-aligned quality assurance

The European Commission should encourage the integration of GenAI governance into ESG-aligned internal and external quality assurance processes (ENQA et al., 2015). This is one of the most practical implementation routes, because quality assurance already provides established

institutional mechanisms for programme design, assessment, student-centred learning, learning resources, student support, public information, monitoring and continuous improvement.

GenAI should be considered in programme review where it affects learning outcomes, assessment methods, student competences or disciplinary practice. It should be considered in assessment policy where it affects authorship, originality, feedback, grading or evidence of learning. It should be considered in student support where AI chatbots, tutoring systems or administrative assistants are used. It should be considered in public information where students and applicants need to know how AI is used in academic processes.

External quality assurance should be introduced carefully, avoiding rigid compliance burdens. The aim should not be to create a checklist that freezes innovation, but to ensure that institutions can demonstrate that they have appropriate policies, safeguards, monitoring mechanisms and improvement processes for GenAI use in learning and teaching.

Embedding GenAI into quality assurance would help institutions move beyond isolated policy documents. It would make GenAI governance part of ordinary institutional responsibility, ensuring that policies are implemented, reviewed, improved and connected to student experience, teacher support and academic standards.

ADMIT dimensions: Accountability, Transparency. AI Act extension: connects AI governance to the established ESG 2015 quality assurance framework for the European Higher Education Area, ensuring that GenAI is governed through existing institutional quality mechanisms rather than treated as a separate technical issue. For students: student-centred quality assurance that includes GenAI governance as a dimension of programme review and public information. For teachers: involvement in quality assurance processes that address GenAI deployment in teaching and assessment. For institutions: integration of GenAI governance into internal and external quality assurance, including ESG-aligned review, accreditation and public reporting.

4.10 Support a European implementation route based on guidance, templates and shared tools

The Commission should support implementation through practical instruments rather than relying only on high-level principles. These instruments could include sector-specific guidance, model institutional policies, disclosure templates, classification checklists, procurement clauses, student-facing notices, appeal-procedure templates, academic integrity guidance, AI literacy resources and self-assessment tools, building on the operational approach already developed in the ADMIT Guide and Institutional Self-Assessment Tool (Triantafyllopoulos et al., 2025a; Triantafyllopoulos et al., 2025b).

The implementation route should be based on cooperation with higher education institutions, student organisations, quality assurance agencies, data protection experts, AI governance bodies, teachers, researchers and civil society actors. This participatory approach would

help ensure that guidance is realistic, context-sensitive and aligned with the diversity of European higher education.

Participation should be formalised through student and staff representation in GenAI governance committees, procurement pilots, impact assessments, incident reviews and periodic evaluation. Institutions should make particular efforts to include representatives of students with disabilities, multilingual learners, international students and groups likely to be disproportionately affected by automated systems.

The Commission should also support piloting and mutual learning. Institutions should be encouraged to test GenAI governance models, share evidence, compare approaches and refine policies in light of practice. This coordination is particularly important as national implementation initiatives are beginning to translate the AI Act into domestic legal and institutional arrangements, creating the need for a common European reference framework for higher education (Presidency of the Council of Ministers, 2026).

This implementation route should preserve institutional autonomy while providing a shared European baseline. Universities should remain free to adapt GenAI governance to their mission, disciplines, national legal context and pedagogical model. However, all institutions should be expected to meet common minimum expectations concerning transparency, accountability, academic integrity, human oversight, equity, data governance, training, monitoring and quality assurance. Designing this baseline requires explicit attention to the diversity of European higher education delivery models. EU guidance should be applicable across campus-based, blended, hybrid, online, distance and transnational modes of provision. The feasibility and proportionality of particular safeguards - including identity verification, assessment supervision, human review and opt-out alternatives - may vary according to the institution's delivery model, student population and national legal context. Guidance should therefore establish common outcomes and safeguards without assuming a single campus-based operational model.

Taken together, these recommendations define a practical and proportionate European response. They do not seek to prohibit LLMs in higher education or to impose a rigid regulatory model. They aim to ensure that the adoption of GenAI remains legally compliant, ethically responsible, pedagogically meaningful, socially inclusive and institutionally accountable.

Policy Implementation Toolkit. *In order to translate these recommendations into institutional practice, the ADMIT project proposes the following Policy Implementation Toolkit. The toolkit components listed below are derived from the analysis, gap identification and recommendation process developed in this deliverable, and draw on the operational tools already produced within the ADMIT partnership (D5.2, D5.3, D5.4). They are intended for adaptation by higher education institutions, quality assurance agencies and national implementation bodies.*

1. Classification Checklist Template. *A structured procedure, informed by the use-case typology of Table 1 (§2.2) and the ADMIT ethical dimensions (Table 2, §2.5), for classifying each GenAI deployment by intended purpose, user groups, affected persons, data processed, role of AI*

outputs, impact on academic outcomes, human oversight arrangements and appeal mechanisms. The checklist supports both AI Act compliance and institutional accountability.

2. Disclosure Notice Template. A model notice for informing students, teachers and staff whenever GenAI is used in teaching, assessment, support or administrative processes. The notice template covers: what AI system is used, for what purpose, what data are processed, what limitations apply, whether outputs are advisory or decision-shaping, what human oversight is in place, and how to access further information, support or appeal. Adaptable to different contexts (course-level, platform-level, institutional-level).

3. Procurement Clause Model. A set of model clauses for institutional procurement of LLM-based systems and services, covering: model identification and versioning, data processing and retention, security and accessibility requirements, language-performance validation for the languages of instruction and assessment, incident notification, update procedures, subcontractor transparency, intellectual property, GDPR compliance, and the right to reassess educational suitability after significant model or system changes. Aligned with the GPAI provider obligations of Articles 51–55 of the AI Act.

4. Academic Integrity Policy Template. A model institutional policy covering: definition of permitted, restricted and prohibited GenAI uses in coursework, examinations and assessment; disclosure requirements for students and teachers; assessment redesign principles for GenAI-rich environments; fair misconduct investigation procedures with human review; treatment of AI-detection tool outputs as indicative rather than conclusive; appeal mechanisms; and periodic review. Informed by the analysis of §3.2 and the ADMIT Guide (D5.3).

5. Self-Assessment Cycle Guide. A structured self-assessment cycle based on the eight ADMIT ethical dimensions (D5.2) and operationalised through the role-specific questions of the ADMIT Institutional Self-Assessment Tool (D5.4). The cycle covers: initial baseline assessment; identification of strengths, weaknesses and gaps; development of an action plan; implementation; follow-up assessment; and reporting. Designed for integration into existing ESG-aligned internal quality assurance processes.

6. Appeal Procedure Template. A model procedure for students affected by AI-mediated or AI-influenced academic decisions (assessment, misconduct detection, admission, progression, proctoring). The template covers: right to information on the role of AI in the decision; right to human review; timeline and process for appeal; evidence access; decision criteria; and documentation requirements. Aligned with the human oversight and contestability principles of the AI Act and the ADMIT dimension of Teacher and Student Agency and Oversight.

These toolkit components are presented as model instruments for institutional adaptation. They do not prescribe a single governance model, but provide a common European baseline that institutions can adjust to their mission, disciplines, national legal context, delivery model and pedagogical approach. The toolkit supports the transition from policy recommendation to institutional practice and is intended to complement the broader guidance, piloting and mutual learning activities described above.

Each of the six model instruments listed above is provided as an editable, operational template rather than a mere description. To this end, every template specifies: mandatory and optional fields; decision rules and risk thresholds; responsible institutional roles; the evidence or documentation required; approval, escalation and review procedures; version-control and update fields; at least one completed higher education example; and explicit cross-references to the relevant AI Act provisions and ADMIT dimensions. The full templates are provided as editable supplementary files in Annex A.

To make the recommendations directly usable by policymakers, the implementation matrix below specifies, for each recommendation, the proposed lead actor, supporting actors, policy instrument, time horizon, expected output and monitoring indicator, distinguishing action at EU, EHEA, national and institutional levels.

Table 5. Implementation Matrix for the policy recommendations

Recommendation	Lead actor	Supporting actors	Policy instrument	Time horizon	Expected output	Monitoring indicator
4.1 Establish EU-level guidance for the classification and deployment of LLM-based systems in higher education	European Commission	AI Office; QA agencies; HEIs	Commission guidance + classification template	Short term	Sector classification guidance and template	% HEIs with a documented classification process
4.2 Define a common transparency standard for educational GenAI use	European Commission	HEIs; QA agencies; student bodies	Guidance + model disclosure notices	Short–medium	Common disclosure standard	% deployments with disclosure notices
4.3 Establish minimum safeguards for AI-mediated academic decisions	EU / national	HEIs; data-protection authorities	Council recommendation	Medium	Minimum safeguards baseline (human review, appeal)	Appeals volume and outcomes
4.4 Develop common European reference principles for institutional academic integrity policies on GenAI	EHEA / HEIs	ENQA; Bologna Follow-Up Group	ESG-aligned guidance	Medium	Model academic-integrity policy	% programmes with GenAI integrity rules
4.5 Ensure equitable access, accessibility and opt-out safeguards	HEIs	Commission; disability organisations	Institutional policy + procurement	Short–medium	Access and accessibility standards	Accessibility conformance; opt-out uptake
4.6 Ensure environmentally proportionate GenAI deployment	HEIs	Commission; suppliers	Procurement + self-assessment	Medium	Environmental proportionality checks	Resource-use indicators in procurement
4.7 Make structured AI literacy and teacher training a precondition for formal institutional deployment	HEIs / national	Commission; Erasmus+	Training programmes; pilots	Short–medium	Literacy and training programmes	Staff/student training coverage
4.8 Introduce periodic institutional self-assessment and EU-level monitoring	HEIs / EU	QA agencies; observatories	Self-Assessment Tool (D5.4); monitoring	Recurring	Periodic self-assessment cycles	Completion; incidents; outcome indicators

4.9 Embed GenAI governance into ESG-aligned quality assurance	QA agencies	ENQA; HEIs	ESG-aligned QA	Medium	GenAI embedded in QA review	Coverage in internal/external QA
4.10 Support a European implementation route based on guidance, templates and shared tools	European Commission	HEIs; AI Office; civil society	Guidance; pilots; shared toolkit	Ongoing	Shared toolkit + mutual learning	Adoption of toolkit instruments

References

- Carbonel, H., Triantafyllopoulos, L., Kalles, D., Kameas, A., and Jullien, J.-M. (2024). Taxonomy for the ethical use of Generative AI in Higher Education. Research Report no. 5.2. Zenodo. doi: 10.5281/zenodo.17201673
- ENQA, ESU, EUA, and EURASHE. (2015). Standards and Guidelines for Quality Assurance in the European Higher Education Area (ESG). Brussels, Belgium.
- European Commission. (2020). Digital Education Action Plan 2021–2027: Resetting education and training for the digital age. Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions, COM(2020) 624 final.
- European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Official Journal of the European Union.
- European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). Official Journal of the European Union.
- Presidency of the Council of Ministers. (2026). Comunicato stampa del Consiglio dei Ministri n. 177. 10 June 2026.
- Triantafyllopoulos, L., Panagiotopoulos, A., Takouli, V., Kalles, D., Kameas, A., and Antonaci, A. (2025a). Institutional Self-Assessment Tool. ADMIT Deliverable D5.4. Zenodo. doi: 10.5281/zenodo.17864328
- Triantafyllopoulos, L., Takouli, V., Kalles, D., Kameas, A., and Antonaci, A. (2025b). Guide of ethical use of LLM / Generative AI systems in Higher Education. ADMIT Deliverable D5.3. Zenodo. doi: 10.5281/zenodo.1786430

Annex A- Policy Implementation Toolkit: Operational Templates

This annex provides the editable, operational versions of the six model instruments introduced in the Policy Implementation Toolkit (§4.9). Each template is designed to be operational rather than merely descriptive: it specifies its purpose and trigger, the responsible institutional roles, mandatory and optional fields, decision rules and thresholds, the evidence required, approval/escalation/review procedures, version-control fields, explicit cross-references to the relevant AI Act provisions and ADMIT dimensions, and at least one completed higher education example. Institutions may adopt the templates as they stand or adapt them to their mission, disciplines, national legal context and delivery model. The templates should be read together with the Implementation Matrix (Table 5, §4.10), which identifies the lead actors, instruments and monitoring indicators for each recommendation.

A.1 Classification Checklist Template

Purpose	Classify each GenAI deployment by intended purpose and risk in order to determine the applicable AI Act obligations and institutional safeguards.
Trigger / when to complete	Before any new GenAI system or use is approved; on material change of model, purpose or user group; and at periodic review.
Responsible roles	Deploying unit (owner); institutional AI governance officer (review); Data Protection Officer (data aspects).
Mandatory fields	System/model name and version; intended purpose; user groups; affected persons; categories of personal data; role of AI outputs (advisory / decision-shaping / decisive); impact on academic outcomes; human-oversight arrangement; appeal mechanism; provisional risk category.
Optional fields	GPAI provider; contract reference; linked DPIA/FRIA reference; environmental resource note.
Decision rules & thresholds	If outputs solely determine an academic decision or affect rights, treat as high-impact and require human review plus FRIA screening; if used for assessment or misconduct detection, require mandatory human confirmation; if purely informational, apply standard safeguards.
Evidence required	Completed checklist; system specification/screenshots; description of the oversight workflow.
Approval, escalation & review	Sign-off by the AI governance officer; escalation of high-impact cases to the institutional AI committee; review at least annually or on change.
Version control	Checklist version, date, author and decision log.
AI Act references	Art. 6 and Annex III; Art. 50 (transparency); Art. 26 (deployer duties); Art. 27 (FRIA screening).
ADMIT dimensions	Accountability; Transparency; Technical Robustness and Safety.
Worked example (higher education)	Automated first-marking of short-answer quizzes is decision-shaping and therefore high-impact: human confirmation is required and FRIA screening is triggered, with appeals handled through Template A.6.

A.2 Disclosure Notice Template

Purpose	Inform students, teachers and staff whenever GenAI is used in teaching, assessment, support or administrative processes.
----------------	--

Trigger / when to complete	At the point of use, enrolment or assessment briefing, and whenever a new GenAI-supported process goes live.
Responsible roles	Course or service owner (drafting); communications/registry (publication).
Mandatory fields	System used; purpose; data processed; known limitations and error modes; whether outputs are advisory or decision-shaping; human oversight in place; how to obtain information, support or appeal; opt-out or alternative where applicable.
Optional fields	Language-availability note; accessibility statement; contact point; version date.
Decision rules & thresholds	Decision-shaping uses require an explicit statement of human oversight and an appeal route; language-sensitive uses require a validity and limitations note (see Template A.3).
Evidence required	Published notice and a record of where and when it was displayed.
Approval, escalation & review	Approved by the course/service owner; template maintained centrally; reviewed each academic year.
Version control	Notice version and effective date.
AI Act references	Art. 50 and Recital 132 (transparency to users); Art. 26(7) for staff-affecting uses.
ADMIT dimensions	Transparency; Teacher and Student Agency and Oversight.
Worked example (higher education)	A platform-level notice for an AI tutoring chatbot states that it is advisory only, has no grading effect, has known multilingual limitations, and offers a human-tutor fallback.

A.3 Procurement Clause Model

Purpose	Provide model contract clauses for the institutional procurement of LLM-based systems and services.
Trigger / when to complete	Any procurement, renewal or major upgrade of a GenAI system or service.
Responsible roles	Procurement office (lead); legal; DPO; IT security; accessibility officer.
Mandatory fields	Model identification and versioning; data processing/retention and sub-processors; security and accessibility (EN 301 549); language-performance validation for the languages of instruction and assessment; incident notification; update/change procedures; subcontractor transparency; intellectual property; GDPR compliance and lawful basis; right to re-assess educational suitability after significant model or system change.
Optional fields	Environmental transparency (resource use); exit and data portability; audit rights; service levels.
Decision rules & thresholds	For high-impact or large-scale services, require documented language validation, accessibility conformance and environmental disclosure before award.
Evidence required	Vendor responses; conformance evidence; validation and test reports.
Approval, escalation & review	Award sign-off by procurement, legal and DPO; re-assessment on major model change.
Version control	Clause-set version aligned to AI Act updates.
AI Act references	Arts. 51-55 (GPAI provider obligations); Art. 25 (responsibilities along the value chain); Art. 50.
ADMIT dimensions	Accountability; Technical Robustness and Safety; Societal, Individual and Environmental Wellbeing.
Worked example (higher education)	A tender for an institution-wide writing assistant requires EN 301 549 conformance, validation in the two languages of instruction, and a resource-use datasheet.

A.4 Academic Integrity Policy Template

Purpose	Provide a model institutional policy on permitted, restricted and prohibited GenAI use in coursework, examinations and assessment.
Trigger / when to complete	Adoption or revision of academic regulations; introduction of new assessment types or new GenAI tools in curricula.
Responsible roles	Academic board (owner); teaching committees; student representatives; registry.
Mandatory fields	Definitions of permitted, restricted and prohibited use; disclosure requirements for students and staff; assessment-redesign principles for GenAI-rich contexts; a fair investigation procedure with human review; treatment of AI-detection outputs as indicative only; appeal mechanism; periodic review.
Optional fields	Discipline-specific annexes; exemplar disclosure statements; staff guidance.
Decision rules & thresholds	AI-detection scores may never be the sole basis for a misconduct finding; every finding requires human review and a right of appeal (Template A.6).
Evidence required	Published policy; assessment briefs; investigation records.
Approval, escalation & review	Academic board approval; annual review; consultation with the student body.

Version control	Policy version and effective academic year.
AI Act references	Art. 50; Art. 14 (human oversight, by analogy); Art. 26(7).
ADMIT dimensions	Academic Integrity and the Safeguarding of Learning; Teacher and Student Agency and Oversight; Fairness.
Worked example (higher education)	A programme permits GenAI for drafting with disclosure, prohibits it in the final oral defence, and redesigns one written exam into a supervised task.

A.5 Self-Assessment Cycle Guide

Purpose	Provide a structured institutional self-assessment across the eight ADMIT ethical dimensions, integrated into ESG-aligned quality assurance.
Trigger / when to complete	The annual quality-assurance cycle, and before and after major GenAI deployments.
Responsible roles	Quality assurance office (lead); AI governance committee; faculty representatives.
Mandatory fields	Baseline assessment; strengths, weaknesses and gaps; action plan with owners and deadlines; implementation record; follow-up assessment; reporting to QA governance; outcome-based indicators for high-impact uses.
Optional fields	Benchmarking; external peer review; linkage to the institutional risk register.
Decision rules & thresholds	High-impact deployments require outcome-based monitoring (error and false-positive rates, appeals, disaggregated effects, accessibility incidents, learning-outcome and staff-workload effects, data-protection incidents and resource use) with defined baselines, thresholds and escalation criteria.
Evidence required	Completed self-assessment; action-plan tracker; monitoring data.
Approval, escalation & review	Reviewed by QA governance; unresolved high-impact risks escalated to senior management; annual cycle.
Version control	Cycle year and tool version (ADMIT D5.4).
AI Act references	Art. 26 (deployer duties); Art. 72 (post-market monitoring, by analogy).
ADMIT dimensions	All eight dimensions, with emphasis on Accountability, Fairness and Wellbeing.
Worked example (higher education)	An annual review flags rising false-positives in an AI proctoring tool: the defined threshold is breached, triggering escalation and external evaluation.

A.6 Appeal Procedure Template

Purpose	Provide a model procedure for students affected by AI-mediated or AI-influenced academic decisions.
Trigger / when to complete	Any assessment, misconduct, admission, progression or proctoring decision in which AI played a role.
Responsible roles	Independent reviewer or panel (not the original decision-maker); registry; student adviser.
Mandatory fields	Right to information on the role of AI in the decision; right to human review; appeal timeline and process; evidence access; decision criteria; documentation and outcome record.
Optional fields	Interim measures; dedicated support for students with disabilities and multilingual learners; representation.
Decision rules & thresholds	Human review is mandatory; the reviewer may overturn or remit; AI outputs are evidence and never the sole determinant.
Evidence required	Original decision record; statement of the AI role; appeal file and outcome.
Approval, escalation & review	Overseen by the academic registrar; escalation to ombuds or senate; periodic review of appeal outcomes.
Version control	Procedure version and effective date.
AI Act references	Art. 14 (human oversight); Art. 50; provisions on redress.
ADMIT dimensions	Teacher and Student Agency and Oversight; Fairness; Accountability.
Worked example (higher education)	A student contests an AI-flagged misconduct case; an independent panel finds the detector unreliable for non-native writing and dismisses the case.

License used: This work is licensed under a Creative Commons Attribution Share Alike 4.0 International License: <https://creativecommons.org/licenses/by-sa/4.0/>

With this license, you are free to **share** copy and redistribute the material in any medium or format. You can also **adapt** remix, transform and **build upon** the material for any purpose, even commercially.

But only Under the following terms:

Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use. **ShareAlike** — If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

Disclaimer: The European Commission support for the production of this publication does not constitute an endorsement of the contents which reflects the views only of the authors, and the Commission cannot be held responsible for any use which may be made of the information contained therein.



**Co-funded by
the European Union**