POLITECNICO DI TORINO
Repository ISTITUZIONALE

Mapping the Policy Landscape of AI in Higher Education: A Multi-level Analysis Across European Institutions

Original
Mapping the Policy Landscape of AI in Higher Education: A Multi-level Analysis Across European Institutions / Zammit, J., Grech, V., Francalanza, E., Maffei, A., Sala, R., Pirola, F., Pezzotta, G., Ljubi, S., Batisti, L., Podržaj, P., Reznichenko, I., Priarone, P.C., Antonelli, D., Pasko, L., Olech, M.. - STAMPA. - 813:(2027), pp. 240-254. (45th IFIP WG 5.7 International Conference on Advances in Production Management Systems, APMS 2026 Bergamo, Italy 13-17 Settembre 2026) [10.1007/978-3-032-38618-2_16].

Availability:
This version is available at: 11583/3017162 since: 2026-10-28T16:19:35Z

Publisher:
Springer Science and Business Media Deutschland GmbH

Published
DOI:10.1007/978-3-032-38618-2_16

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

28 November 2026

Mapping the Policy Landscape of AI in Higher Education: A Multi-level Analysis Across European Institutions

Joseph Zammit[1], Victoria Grech[1], Emmanuel Francalanza[1], Antonio Maffei[2], Roberto Sala[3](✉), Fabiana Pirola[3], Giuditta Pezzotta[3], Sandi Ljubić[4], Luka Batistić[4], Primož Podržaj[5], Igor Reznichenko[5], Paolo C. Priarone[6], Dario Antonelli[6], Lukasz Pasko[7], and Marcin Olech[7]

[1] University of Malta, Faculty of Engineering, Msida, Malta
Joseph.Zammit@um.edu.mt
[2] KTH Royal Institute of Technology, Stockholm, Sweden
[3] University of Bergamo, Dalmine, Italy
roberto.sala@unibg.it
[4] University of Rijeka, Faculty of Engineering, Rijeka, Croatia
[5] University of Ljubljana, Ljubljana, Slovenia
[6] Politecnico di Torino, Torino, Italy
[7] Rzeszów University of Technology, Rzeszów, Poland

Abstract. Generative AI has turned AI governance in higher education from a strategic concern into an operational issue affecting assessment, academic integrity, data protection, procurement, and institutional accountability. Yet comparative analysis across European, national, and institutional policy levels remains limited. This paper presents a multi-level comparative document analysis of 56 policy and guidance documents across seven partner cases in six European countries. Using a common review protocol and an eight-dimensional coding framework, it examines how ethics, privacy, inclusion, transparency, governance, technical standards, accountability, and sustainability are articulated across policy levels. The findings show a clear maturity gradient: European instruments provide strong normative and legal anchors, national documents remain largely strategic, and institutional policies are the most operational but also the narrowest in scope. Strongest coverage concerns ethics, privacy, transparency, and academic integrity, while governance, ownership, liability, interoperability, procurement, and sustainability remain weak. A recurring pedagogical integration gap is also evident. The study, therefore, provides a comparative policy map of AI governance in European higher education and highlights where multi-level misalignment continues to limit coherent institutional operationalisation.

Keywords: Higher Education · Generative AI · Policy Analysis · Digital Governance · Academic Integrity.

© The Author(s) 2027
Published by Springer Nature Switzerland AG 2026
G. Pezzotta et al. (Eds.): APMS 2026, IFIP AICT 813, pp. 240–254, 2027.
https://doi.org/10.1007/978-3-032-38618-2_16

1 Introduction

Generative AI has become a live governance issue for higher education institutions (HEIs) because it unsettles long-standing arrangements around authorship, assessment, transparency, data protection and institutional responsibility. The rapid mainstreaming of tools such as ChatGPT has intensified the need for universities to define acceptable use, protect sensitive data, redesign assessment, and clarify the relationship between human judgment and machine assistance [1, 4, 6]. These pressures extend beyond teaching and learning to research workflows, administrative processes, procurement decisions, and the governance of external AI services [1, 2, 5].

In Europe, these questions are shaped by a layered policy environment. The AI Act establishes a risk-based legal framework for safe and trustworthy AI [1]. GDPR continues to provide the central legal architecture for lawful, fair, transparent, and secure processing of personal data [2]. The Digital Education Action Plan positions digital readiness, educator competence, and inclusive infrastructure as strategic priorities for European education systems [3]. UNESCO and European Commission guidance further reinforce the importance of human-centred, ethically governed, and pedagogically meaningful AI adoption in education [4]–[6]. Together, these instruments establish a strong normative and legal baseline. What remains less clear is how far this baseline translates into institutionally usable guidance on disclosure, assessment design, tool approval, data entry into external systems, and responsibility allocation when AI-related risks or incidents arise.

Recent literature shows both the pace of institutional response and its unevenness. AI research in higher education (HE) has grown rapidly, but important ethical, methodological, and contextual gaps remain [7]. Studies of university GenAI policies in the United States and Asia find that institutions are often open but cautious, typically prioritising academic integrity, privacy, and guidance for staff and students while leaving deeper governance questions only partially addressed [8]–[10]. Humble's analysis of Swedish university policies argues that AI governance in HE requires alignment across strategic, pedagogical, ethical-legal, operational, and adaptive domains [11]. At the same time, academic-integrity scholarship has shown that reactive prohibition is an insufficient response, especially where AI has become embedded in ordinary learning practices [9, 12, 13].

Important as these contributions are, the research base still has a clear gap. Much of the literature focuses on AI in teaching and learning, AI ethics, academic integrity, or institutional policy within a single national setting. Comparatively underexplored is how policy intent travels across governance levels in a multi-country European context. The central question is therefore not only what universities say about AI, but how European legal and policy expectations are translated through national frameworks and then operationalised within institutions.

This paper addresses that gap through a comparative document analysis of AI-related policy instruments across European, national, and institutional levels. It focuses on three main issues: (i) variation in governance expectations across the three levels, (ii) the extent to which key governance dimensions are consistently covered or left underdeveloped, and (iii) how pedagogical integration is treated within the wider landscape of ethics, integrity, and governance guidance.

2 Background and Related Work

2.1 AI in Teaching and Learning

Much of the literature on AI in HE has focused on applications, opportunities, and learner-facing implications. Bond et al. identify rapidly expanding evidence base but argue that the field still requires greater ethical attention, interdisciplinary collaboration, and methodological rigour [7]. Public debate accelerated sharply after the release of accessible generative tools, pushing universities to respond in real time to technologies that had moved into ordinary educational practice.

At the practice level, Wang et al. show that universities in the United States have generally adopted an open-but-cautious stance toward GenAI, frequently supplementing policy with practical resources such as syllabus templates, staff support, workshops, and explanatory materials [8]. Their findings suggest that governance in HE operates through both formal policy and softer implementation artefacts, including templates, guidance notes, exemplars and pedagogical support. This is also relevant for European HEIs, where institutional readiness may depend as much on operational support materials as on formal policy statements.

2.2 AI Ethics, Integrity and Academic Practice

A second body of literature examines AI through the lens of ethics, integrity, and responsible educational use. UNESCO's guidance frames GenAI adoption as a human-centred governance challenge, stressing transparency, equity, teacher development, and the preservation of human agency in educational decision making [4]. Similarly, the European Commission's ethical guidelines for educators emphasise human oversight, fairness, privacy, technical robustness, accountability, and the need for reflective professional judgement when AI is introduced into teaching and learning [6]. Together, these frameworks help establish the normative vocabulary of responsible AI in education.

Academic-integrity research adds a more practice-specific dimension. Cotton et al. argue that the emergence of ChatGPT and similar tools requires universities to move beyond simplistic cheating narratives and instead rethink assessment design, learning support and integrity cultures more holistically [12]. Moya et al., also show that AI creates both bounded and unbounded ethical implications for HE, with implications for misconduct analysis, legitimate assistance, and equitable decision making [13]. Luo's critical review of GenAI policies in HE assessment deepens this point by showing that many policies are still built around the defence of originality, often without adequately reconsidering what originality and authentic learning should mean in AI-mediated environments [9]. Taken together, these studies show that integrity governance cannot be reduced to prohibition or detection, but must also engage with pedagogy, assessment, disclosure, and the conditions under which students can learn responsibly with AI.

2.3 AI Governance and Policy in Higher Education

A third strand of research addresses institutional governance and policy more explicitly. Dai et al., in their study of Asian university policies, identify both promising practices

and recurrent gaps, including uneven specificity, variation in disclosure expectations, and limited attention to longer-term governance structures [10]. Humble's study of Swedish university guidelines similarly demonstrates that HE AI policy now requires alignment across multiple institutional domains, including pedagogy, ethics and law, operations, strategy, and adaptive review [11]. These studies move beyond general commentary on AI and toward the organisational question of how universities govern AI institutionally.

This question also intersects with the broader literature on digital governance and digital transformation in HE. Carmo et al. show that institutional digital transformation is not simply a matter of adopting tools, but of building coherent managerial, organisational, and strategic capabilities [14]. Earlier work on IT governance in HE also highlights the importance of clear roles, structures, controls, and decision mechanisms in contexts where technology becomes deeply embedded in institutional functioning [15]. AI governance should therefore be understood within the wider context of digital governance, since universities must not only decide whether AI is educationally useful, but also whether they possess the governance arrangements needed to adopt, oversee and revise its use responsibly over time.

2.4 European Policy Baseline

The European policy environment gives this study its multi-level significance. The AI Act is the most visible development, establishing a risk-based regulatory model [1]. Even where many universities fall outside the strictest high-risk categories, the Act still reshapes expectations around transparency, AI literacy, accountability, and risk management. GDPR remains equally central because AI use in HE often involves personal data, behavioural data, assessment data, research data or sensitive information that may be processed through external platforms [2]. Universities, therefore, remain subject to substantial obligations regarding lawful processing, minimisation, confidentiality, and governance over third-party services.

The Digital Education Action Plan adds a broader systemic dimension. It frames digital transformation in education through two linked priorities: the deployment of digital technologies and the development of digital competences [3]. The European Digital Education Hub's AI report reinforces this by connecting governance, teacher competences, assessment, institutional support, and the risks and opportunities of AI in education [5]. Taken together, these European frameworks provide strong legal, strategic, and normative direction, but they remain less specific regarding institutional workflows for assessment, procurement, staff development, and governance ownership.

2.5 Research Gap

The reviewed literature shows that HE policy around AI is expanding quickly, but unevenly. There is already useful work on the ethics of AI, academic integrity and institutional policies in specific national or regional contexts. Less developed, however, is a structured comparative analysis of how policy intent translates across European, national, and institutional levels. This matters because HEIs do not operate in a single policy environment but interpret and operationalise layered frameworks that differ in legal force, specificity, timing, and educational relevance.

This paper addresses a narrower, more precise problem than the broader question of AI in education. It examines how governance intent moves across policy levels, how it is operationalised at different stages, and how institutions respond to variation in national or sectoral guidance. This focus moves the analysis beyond generic claims about ethics or innovation and toward a more grounded account of governance maturity, implementation asymmetry, and multi-level policy coordination in European HE.

3 Methodology

3.1 Study Design and Analytical Scope

The study adopted a multi-national, desk-based policy analysis designed to map the policy environment shaping responsible AI use in HE. Rather than compiling a complete legal inventory for each jurisdiction, the study focused on the most policy-salient instruments likely to affect HEI teaching, assessment, research, administration, and governance. The design was therefore purposive and comparative: purposive in document selection, and comparative across governance levels and case contexts.

The final corpus comprised 56 unique documents distributed across European ($n = 7$), national ($n = 20$), and institutional ($n = 29$) levels. These included laws, regulations, strategic plans, digital education strategies, research integrity guidance, conduct codes, institutional AI guidelines, and related governance texts. The analysis covered six national settings and seven comparative case units in total, including two Italian institutional contexts. The reviewed material spans from 2012 to 2025, thereby capturing both pre-GenAI baseline governance instruments and the more recent wave of post-2022 AI-specific policies.

3.2 Inclusion and Exclusion Logic

Documents were included where they satisfied two conditions. First, they had to be publicly attributable and institutionally or governmentally authoritative. Second, they had to be materially relevant to AI governance in HE, either directly or indirectly. Directly relevant documents explicitly addressed AI, generative AI, or responsible use of AI tools. Indirectly relevant documents were included where they set conditions that meaningfully shape institutional AI adoption, such as data-protection rules, research-integrity guidance, digital-transformation strategies, digital-education plans, student conduct rules or institutional strategic plans.

This broader inclusion logic was necessary because AI governance in HEIs is rarely contained within documents explicitly labelled as AI policy. In many settings, operational governance is distributed across adjacent instruments, including data-handling rules, staff conduct codes, and research-integrity guidance. Documents were excluded if they were primarily promotional, journalistic, vendor-produced, or lacking normative content. Purely technical product documentation was excluded, as were duplicate versions where a later or consolidated version was available. Internal analytic templates and unpublished partner reports were treated as part of the research process rather than as formal references in the paper.

3.3 Multilingual Review Protocol and Cross-Case Standardisation

A common structured review template was used across all cases to standardise document metadata, thematic extraction and interpretive notes. For each selected document, reviewers recorded title, issuing authority, year, type, scope, and short descriptive metadata. They then synthesised document content across eight common governance categories.

Because the corpus spans multiple languages, multilingual handling was a central methodological issue. National and institutional documents were reviewed in their original languages by the relevant country partners, while the common analytic template and synthesis outputs were prepared in English. This approach preserved contextual and linguistic familiarity at the point of interpretation while still enabling cross-case comparison through a shared reporting language. Where exact wording was important, key extracts were translated into English and accompanied by interpretive notes to preserve context.

Not all cases reviewed the same number of documents. The case sets ranged from approximately six to eleven documents, reflecting the uneven maturity of policy landscapes, document availability, and the relative presence or absence of AI-specific institutional guidance at the time of review. This variation was analytically meaningful, since smaller document sets may themselves indicate thinner governance environments.

3.4 Coding Framework and Consistency Procedures

The analysis used a directed qualitative content-analysis approach built around eight pre-defined governance dimensions:

- Ethics and Integrity
- Data Privacy and Security
- Equity and Inclusion
- Transparency and Explainability
- Governance Structures
- Technical Standards and Interoperability
- Accountability and Liability
- Sustainability

These dimensions were chosen because they capture the most recurrent governance questions raised by public AI frameworks and by the emerging HE literature [1, 4]–[6, 11]. For each document and category, reviewers extracted evidence and added interpretation notes on coverage, strengths, weaknesses and implementation implications. The coding structure, therefore, combined source traceability with analytic judgement.

Coding consistency was supported by a shared template, a quote-plus-interpretation structure, and a central harmonisation stage that checked metadata, scope labels, and category use across cases. Consistency was pursued through structured guidance, iterative clarification, and auditable pairing of extracts with interpretations.

3.5 Synthesis Logic

The synthesis proceeded in two layers. The first was within-case synthesis, where each partner case was treated as a coherent national-institutional policy environment. The

second was cross-case synthesis, where evidence was compared across governance levels and across cases to identify recurring patterns, asymmetries and omissions.

To support comparative presentation, the study also developed a three-point interpretive rating for selected dimensions, namely national operational support, institutional operational guidance, and pedagogical integration support. These ratings were assigned based on convergent documentary evidence and cross-case analytic notes and should be read as comparative heuristics rather than measured outcomes.

3.6 Methodological Limitations

Documentary analysis is well-suited to studying policy intent, but it cannot, by itself, establish the quality of actual implementation or everyday institutional practice. A university may possess a strong formal policy that is poorly implemented or, conversely, govern effectively through informal support practices that are only partially documented. Therefore, the method captures stated governance environments rather than lived organisational behaviour.

A second limitation is corpus heterogeneity. Laws, strategic plans, conduct codes and institutional guidance differ in legal force, specificity, intended audience, and temporal horizon. This makes a strict like-for-like comparison impossible. In addition, national and institutional policy landscapes vary in maturity, affecting the density of observable governance signals. A further limitation concerns external transferability: because the study examines a specifically European multi-level governance environment shaped by EU, national, and institutional interactions, the findings should not be assumed to generalise directly to regions with different regulatory architectures or more centralised higher-education governance systems.

A further limitation concerns temporality. Several strategic documents in the corpus predate the rapid post-2022 diffusion of generative AI. Accordingly, the absence of explicit reference to GenAI should not necessarily be interpreted as evidence of neglect but may instead reflect the timing of document adoption. Finally, although the structured template increases consistency, interpretive judgement remains intrinsic to qualitative documentary analysis. The results should therefore be read as a robust comparative interpretation rather than a claim of exhaustive legal completeness.

4 Results

4.1 Document Profile and Policy-Level Effects

The first finding concerns the corpus's shape. Of the 56 unique documents reviewed, 7 were European-level instruments, 20 were national, and 29 were institutional, as illustrated in Fig. 1. Therefore, the dataset is weighted toward the institutional level, which is analytically important because institutional policies are where abstract legal and strategic expectations most often become operational. This institutional predominance does not diminish the importance of European and national levels; rather, it suggests that HEI governance is increasingly being operationalised downstream of higher-order frameworks.

The coding distribution reinforces this pattern. After excluding blank placeholders, the evidence base comprised 448 populated coded entries, of which 232 were institutional, 158 national and 58 European as illustrated in Fig. 1. Therefore, Institutional documents generated just over half of the coded evidence, followed by national documents, with European documents contributing the smallest share. This does not imply low importance at the European level. Rather, it indicates that operational detail tends to accumulate downstream, while European documents provide the normative and legal baseline from which later translation occurs.

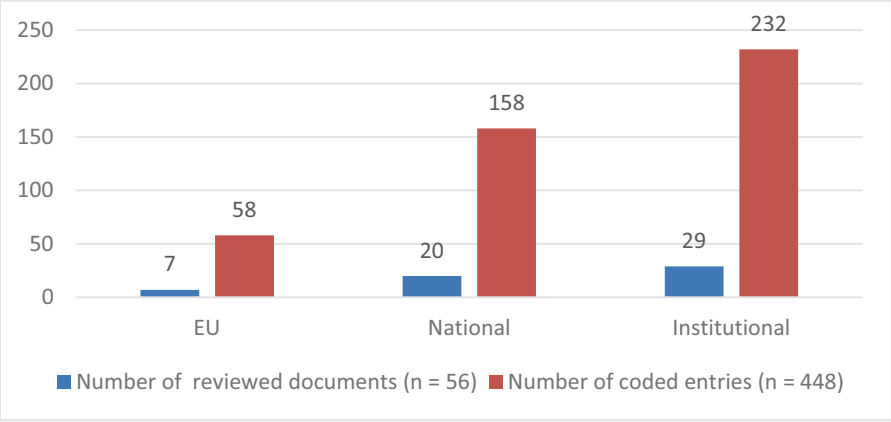

Fig. 1. Policy documents and coded entries by level

The dataset also spans a long-time window, from 2012 to 2025, but the AI-specific part of the data is heavily concentrated in the period from 2023 to 2025. Earlier documents tend to be conduct codes, data-protection instruments, or strategic plans that became newly relevant once AI moved into mainstream academic practice. This temporal pattern points to a layered policy landscape in which earlier governance documents provide baseline controls, while more recent AI-specific guidance seeks to address the operational challenges created by the generative AI uptake.

4.2 A Maturity Gradient Across Governance Levelsd

The clearest overall result is a maturity gradient across the three governance levels. European documents are strongest in their legal and normative framing. They articulate human-centred AI, risk-based governance, data protection, digital competence, fairness, transparency, and accountability [1]–[6]. However, they do not generally function as HEI-ready toolkits. They rarely specify, for example, how acceptable-use expectations should be operationalised at the programme or departmental level, or how the internal approval and oversight process for AI tools should be structured.

National-level documents commonly mediate European intent, but unevenly. In most cases, they echo European values and strategic priorities like innovation, public-sector modernisation, competitiveness, digital readiness, and trustworthy AI. However, they

remain strategic rather than procedural. In some countries, HE appears only indirectly through wider digitalisation or education strategies. In others, AI is addressed nationally, but not with sufficient specificity for HEI operational use. This leaves universities to resolve concrete issues locally.

Therefore, institutional documents are the most operational level in the corpus. They are where disclosure expectations, staff guidance, data-input restrictions, misconduct rules and assessment-related positions are most visible. Yet institutional maturity is uneven, and the strongest institutional policies remain narrower than the governance challenge they aim to address. In many cases, institutions govern AI most clearly in teaching and assessment, but less so in procurement, enterprise risk ownership, technical integration, or non-academic uses.

4.3 Comparative Case Patterns

To support cross-case comparison, the analysis examined three dimensions: national operational support, institutional operational guidance, and pedagogical integration support. These dimensions were used as interpretative comparison devices rather than as a formal index.

Across the seven case contexts, three broad patterns emerged. First, several cases showed limited HE-specific operational detail at the national level, while institutional guidance translated broad principles into more usable practices through staff guidance, course-level statements, disclosure expectations, and assessment-related advice. Second, some cases displayed a layered governance pattern in which newer GenAI guidance sat alongside older strategic documents not originally designed around AI. This should not be interpreted as an absence of governance, but rather as uneven temporal adaptation across policy layers. Third, several cases showed comparatively strong treatment of ethics, privacy, transparency, and responsible-use expectations, while remaining less developed in pedagogical methodology, technical enablement, procurement oversight, or enterprise-level governance arrangements. Taken together, the cross-case pattern suggests that institutions often compensate locally where national operational guidance remains broad, but that institutional proactivity does not automatically resolve pedagogical integration, infrastructure readiness, or long-term governance ownership (Fig. 2).

4.4 The Strongly Developed Governance Dimensions

Across the reviewed corpus, four governance areas show the strongest and most recurrent coverage. The first is ethics and integrity. Nearly every policy environment invokes responsible use, human oversight, fairness or academic honesty. The second is privacy and data security. This is especially visible where institutions warn staff and students not to enter personal, confidential, or sensitive information into external AI systems. The third is transparency, most often in the form of disclosure expectations, traceability requirements, or explicit mention that AI assistance should be acknowledged. The fourth is academic-integrity governance in teaching and assessment, where institutions increasingly define acceptable and unacceptable uses of AI.

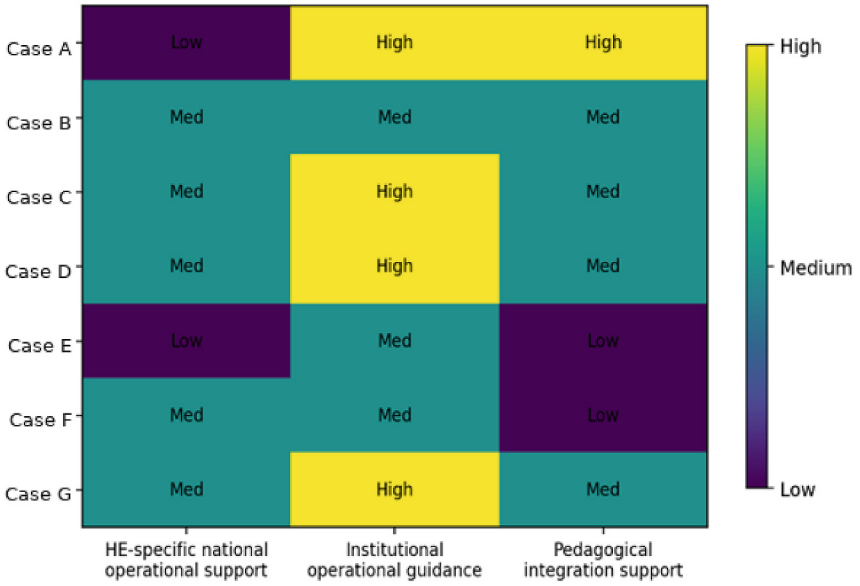

Fig. 2. Comparative governance maturity matrix

One reason these areas are comparatively strong is that they can be translated into behavioural or procedural rules. "Disclose AI use," "verify the output," and "do not upload sensitive data" are relatively actionable norms. They are easier to draft, communicate, and enforce than governance provisions requiring complex organisational design or technical assurance.

4.5 The Weakest Governance Dimensions

The weakest and most unevenly governed dimensions are governance structures, accountability and liability, technical standards and interoperability, procurement oversight and sustainability. Policies frequently state that responsibility remains with human actors, but often do not specify how it is allocated across the AI adoption lifecycle. In many cases, documents do not clearly define decision rights for tool approval, ownership of monitoring responsibilities, incident handling, or authority over AI-enabled processes in student support, administration, or research.

Technical and interoperability issues are even more weakly governed. Many documents offer limited detail on integration with learning management systems, documentation standards, enterprise architecture, vendor monitoring, or update governance. In practice, this means that institutions often govern the visible front-end AI use more clearly than the technical and organisational systems that sustain it.

Sustainability is present in the corpus but usually only in passing. Some documents explicitly acknowledge the environmental footprint of AI, the energy demands of large models, or broader green transition goals. Yet sustainability rarely becomes an operational AI governance requirement. It is more often a contextual value than a decision criterion.

4.6 The Pedagogical Integration Gap

One of the most important findings of the study is that pedagogical integration remains weakly governed across much of the policy landscape. This does not mean pedagogy is absent altogether. Rather, it means that policy attention is disproportionately concentrated on whether AI use is permitted, how it should be disclosed, and how misuse should be limited, while much less attention is paid to how AI should be integrated into teaching and learning in a pedagogically coherent way.

This gap has several dimensions. Firstly, many documents say little about curriculum design. They rarely explain how AI should be aligned with course outcomes, disciplinary purposes, or developmental stages of learning. Secondly, assessment design is often addressed defensively rather than constructively. Institutions warn against misuse, call for disclosure, or encourage redesign of assessments, but often stop short of explaining in detail what robust redesign looks like. Moreover, there is limited treatment of how teachers and students should share responsibility in AI-supported learning processes, especially where AI use may be legitimate but pedagogically variable.

The documentary evidence suggests that this gap is more visible in some cases than others, particularly where ethics, integrity, and governance language have developed faster than pedagogical methodology. In fewer cases, pedagogical operationalisation is more visible through templates, course-level statements, or differentiated rules for use. Overall, however, the field remains more advanced in governance-of-risk than in governance-for-learning.

This matters because pedagogy is where governance ultimately meets educational practice. A university can have strong disclosure rules and still offer weak guidance on how to redesign authentic tasks, how to scaffold responsible AI use, or how to protect learning outcomes that depend on individual judgment, reflection, or iterative problem-solving. If policy stops at control and compliance, it risks governing AI as an integrity problem rather than as a learning-design challenge.

4.7 Illustrative Governance Patterns

Selected examples from the corpus reinforce the broader patterns identified above. Operational maturity appears strongest where institutions govern through usable artefacts, such as templates, course-level statements, disclosure expectations, or differentiated guidance on permitted and prohibited uses. By contrast, some documents show that protective compliance around privacy and confidentiality is more developed than enabling infrastructure or secure institutional toolchains. Other examples indicate that transparency becomes operational when disclosure and traceability are treated as routine workflow requirements rather than abstract principles. The evidence also suggests that risk-control measures do not automatically amount to pedagogical integration, and that sustainability is recognised more often rhetorically than operationally. Finally, several cases illustrate a broader pattern of national–institutional asymmetry, in which national ambition remains strategic while institutions supply the operational detail that educators and staff actually need.

5 Discussion

The comparative analysis suggests that European AI governance in HE is ambitious but uneven in execution. European-level policy does several things well. It provides a common normative and legal vocabulary. It makes risk, rights, privacy, accountability and digital competence difficult to ignore. It also provides institutions with a legitimate basis for justifying local governance decisions. In this sense, the European level is not weak; it is foundational.

The findings suggest that AI governance in HE is best understood as a problem of translation rather than policy absence. This argument concerns policy intent and governance design rather than measured implementation performance at institutional level. European-level instruments provide a strong legal and normative baseline, while national frameworks often translate that baseline into strategic ambition without producing HEI-ready operational mechanisms. Institutions are therefore left to resolve practical questions locally. This does not make institutions weak endpoints in the governance chain. On the contrary, the analysis shows that they frequently act as active translators, compensating for upstream vagueness by developing locally actionable guidance for teaching, assessment, disclosure, and data use.

However, this compensatory process comes with a structural trade-off. As governance becomes more operational downstream, it often becomes narrower in scope. Institutional policies are typically strongest where risks are immediate and visible, especially in assessment, privacy, disclosure, and academic integrity, but weaker in procurement, technical integration, enterprise ownership, interoperability, and longer-term governance arrangements. The result is not simply uneven maturity, but an operational narrowing effect: governance gains procedural depth while losing breadth. This helps explain why universities can appear simultaneously proactive and underprepared.

The recurring pedagogical integration gap should also be understood as a governance-design problem, not only as a pedagogical one. Governance regimes are advancing more quickly in defining whether students may use AI than in addressing how teaching, learning, and assessment should be redesigned in contexts where AI is already present. This has practical consequences. Without stronger pedagogical guidance, institutions risk producing inconsistent student experiences across programmes and modules, placing too much discretion on individual educators, and framing AI primarily as a threat rather than as a pedagogically governable reality. The literature has already suggested that reactive originality-based approaches are insufficient [9, 12], and the documentary evidence here supports that judgement. Policies that only defend against misuse may not help institutions cultivate responsible and educationally purposeful use.

The findings also clarify what a more mature governance architecture would entail. This would require clearer institutional responsibility for approval, monitoring, escalation, and review, as well as governance workflows that go beyond prohibition alone by incorporating approved-tool processes, disclosure norms, and incident pathways. It also requires pedagogical guidance that is more closely connected to curriculum and assessment design, especially in disciplines where AI use is neither fully avoidable nor uniformly appropriate. Stronger specification is also needed for procurement, vendor governance, and data-control arrangements, while sustainability and equity should be operationalised as review criteria rather than treated simply as background values.

The scientific value of the study lies not simply in showing that policies vary, but in showing how they vary across governance levels and where translation weakens. This multi-level perspective helps explain why universities can appear both relatively well governed and structurally underprepared at the same time. They may, for example, possess strong acceptable-use guidance for coursework while still lacking a clear enterprise governance model for AI adoption.

More broadly, this study suggests that governance maturity in HE should be judged not only by the presence of an AI policy, but also by the extent to which that policy connects strategic intent, operational procedures, and pedagogical practice. On that basis, the current European landscape appears to be transitioning from early ethical framing toward more specific institutional operationalisation, but it has not yet achieved stable multi-level alignment. Comparative policy mapping matters because it reveals where policy design remains insufficiently specified, where institutions are compensating for these limitations, and where the next phase of governance work should be concentrated.

6 Conclusion

This study shows that AI governance in European HE is best understood as a problem of translation rather than adoption. While the European policy environment provides a strong and broadly coherent legal and normative foundation, governance intent becomes more uneven as it moves across national and institutional levels. European instruments establish important principles and obligations, national frameworks often translate these into strategic ambitions, and institutions are left to resolve many of the practical questions of implementation for themselves.

The analysis suggests that institutions are not simply passive endpoints in this policy chain. Rather, they frequently compensate for upstream vagueness by producing locally actionable guidance on teaching, assessment, disclosure, acceptable use, and data handling. However, this institutional compensation effect is accompanied by a narrowing of governance scope. As policy becomes more operational, it also tends to become more focused on the most immediate and visible issues, particularly academic integrity, privacy, transparency, and assessment-related practice. Broader questions of procurement, technical integration, enterprise ownership, accountability pathways, interoperability, and sustainability remain less fully developed.

A recurring pedagogical integration gap further reinforces this pattern. Although institutions increasingly regulate whether AI may be used, they less often provide equally developed guidance on how AI should be embedded into curriculum design, assessment strategy, and learning outcomes. This suggests that pedagogy should be treated as a core governance concern rather than as a secondary implementation matter left to individual educators.

The study therefore contributes a comparative policy map of AI governance in European HE, a governance-maturity interpretation across policy levels, and a clearer account of where multi-level misalignment persists. More mature AI governance in HEIs will require not only institutional rules, but better alignment across policy layers, clearer allocation of responsibilities, stronger pedagogical guidance, and governance arrangements that extend beyond behavioural control to include organisational, technical, and strategic

coordination. Future research should therefore examine not only policy texts, but also how these governance arrangements are enacted, contested, and revised in institutional practice.

Acknowledgments. Co-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 Foundation for the Development of the Education System. Neither the European Union nor the Foundation for the Development of the Education System can be held responsible for them. This research was carried out within the framework of the **Erasmus + Funded project PATH-AI** (2025-1-PL01-KA220-HED-000359600).

Disclosure of Interests The authors have no competing interests to declare.

References

1. European Parliament and Council: Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Off. J. Eur. Union. (2024)
2. European Parliament and Council: Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (General Data Protection Regulation). Off. J. Eur. Union. (2016)
3. European Commission: Digital education action plan 2021–2027: resetting education and training for the digital age. Belgium, Brussels (2020)
4. UNESCO: Guidance for generative AI in education and research. UNESCO, Paris, France (2023)
5. European Digital Education Hub: AI report: artificial intelligence in education. European Commission, Brussels, Belgium (2023)
6. European Commission: Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators. Publications Office of the European Union, Luxembourg (2022)
7. Bond, M., et al.: A meta systematic review of artificial intelligence in higher education: a call for increased ethics, collaboration, and rigour. Int. J. Educ. Technol. High. Educ. **21**, 4 (2024). https://doi.org/10.1186/S41239-023-00436-Z
8. Wang, H., Dang, A., Wu, Z., Mac, S.: Generative AI in higher education: seeing ChatGPT through universities' policies, resources, and guidelines. Comput. Educ. Artif. Intell. **7**, 100326 (2024). https://doi.org/10.1016/j.caeai.2024.100326
9. Luo, J.: A critical review of GenAI policies in higher education assessment: a call to reconsider the 'originality' of students' work. Assess. Eval. High. Educ. **49**(5), 651–664 (2024). https://doi.org/10.1080/02602938.2024.2309963
10. Dai, Y., Lai, S., Lim, C.P., Liu, A.: University policies on generative AI in Asia: promising practices, gaps, and future directions. J. Asian Public Policy. **18**(2), 260–281 (2025). https://doi.org/10.1080/17516234.2024.2379070
11. Humble, N.: Higher education AI policies—a document analysis of university guidelines. Eur. J. Educ. **60**, e70214 (2025). https://doi.org/10.1111/ejed.70214
12. Cotton, D.R.E., Cotton, P.A., Shipway, J.R.: Chatting and cheating: ensuring academic integrity in the era of ChatGPT. Innov. Educ. Teach. Int. **61**, 228–239 (2024). https://doi.org/10.1080/14703297.2023.2190148
13. Moya, B.A., et al.: Academic integrity and artificial intelligence in higher education contexts: a rapid scoping review. Can. Perspect. Acad. Integr. **7**(3) (2024). https://doi.org/10.55016/ojs/cpai.v7i3.78123

14. Carmo, J.E.S., Lacerda, D.P., Klingenberg, C.O., Piran, F.A.S.: Digital transformation in the management of higher education institutions. Sustain. Futur. **9**, 100692 (2025). https://doi.org/10.1016/j.sftr.2025.100692
15. Bianchi, I.S., Sousa, R.D.: IT governance mechanisms in higher education. Procedia Comput. Sci. **100**, 941–946 (2016). https://doi.org/10.1016/j.procs.2016.09.253

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

