



Ethical Readiness for Responsible AI in Higher Education

A Pre-deployment Institutional Governance Framework

Mohamed Hussien Abdelaziz

Ejada Systems, Cairo, Egypt

June 2026

Keywords

Artificial intelligence; higher education; AI governance; ethical readiness; institutional accountability; academic integrity

Abstract

The rapid integration of Artificial Intelligence (AI) across higher education institutions has generated substantial governance challenges concerning accountability, transparency, oversight, and responsible deployment. While existing Responsible AI governance frameworks provide normative principles for ethical AI practice, a persistent gap remains between institutional governance commitments and the governance capability required to operationalise those commitments effectively. This paper introduces Ethical Readiness as a governance-precondition framework designed specifically for higher education environments. Ethical Readiness is defined as the institutional governance condition in which accountability structures, governance ownership mechanisms, oversight capability, transparency preparedness, corrective governance capacity, governance adaptability, and governance proportionality are sufficiently developed to support Responsible AI deployment before operational implementation occurs. Adopting a conceptual-theoretical research design, the paper synthesises scholarship across Responsible AI governance, higher education governance, organisational readiness, governance maturity, and institutional legitimacy theory. The framework is operationalised through eight interdependent governance dimensions and a seven-stage governance lifecycle model. Institutional applicability is illustrated through comparative analysis of five AI deployment contexts common in higher education, supported by evidence from emerging university governance initiatives. The study contributes to Responsible AI governance scholarship by introducing governance preparedness as a distinct analytical construct, differentiating Ethical Readiness from governance maturity, organisational readiness, and Responsible AI frameworks, and providing a structured governance architecture capable of supporting institutional reflection and future empirical investigation.

Corresponding Author: Mohamed Hussien Abdelaziz, Ejada Systems, Cairo, Egypt, Email: mh.abdelaziz@ejada.com. Abdelaziz, Mohamed Hussien: "Ethical Readiness for Responsible AI in Higher Education: A Pre-deployment Institutional Governance Framework". *Journal of Ethics in Higher Education* 8(2026): 89–117. DOI: <https://doi.org/10.26034/fr.jehe.2026.9532> © the Author. CC BY-NC-SA 4.0. Visit <https://jehe.globethics.net>

1. Introduction

Artificial intelligence is reshaping higher education at a pace that governance structures have struggled to match. Universities now deploy AI-enabled technologies across teaching, student support, research, assessment, and institutional decision-making, drawn by genuine gains in efficiency, personalisation, and analytical capacity. Yet the same technologies that create educational opportunity also introduce substantial governance challenges around accountability, transparency, fairness, privacy, and academic integrity. These concerns carry particular weight in universities, whose obligations differ meaningfully from those of commercial organisations: institutions entrusted with educational development, knowledge creation, and public accountability face consequences when governance fails that extend well beyond operational inconvenience to matters of institutional legitimacy and educational equity.

The governance response to this challenge has been vigorous at the normative level. International standard-setting bodies have produced influential frameworks — the OECD Artificial Intelligence Principles, the UNESCO Recommendation on the Ethics of Artificial Intelligence, the NIST AI Risk Management Framework, and ISO/IEC 42001 — that collectively establish shared expectations around accountability, human oversight, transparency, and institutional responsibility. Within higher education specifically, a growing number of universities have introduced responsible AI policies, academic integrity guidance, and governance oversight mechanisms.

Despite these developments, a more fundamental governance problem persists. Many institutions publicly endorse responsible AI principles while simultaneously lacking the capability to operationalise them. Policies exist; governance committees are formed; ethical declarations are published. Yet the accountability structures, oversight mechanisms, coordination processes, and corrective capacities that would transform those commitments into effective practice often remain underdeveloped. The gap, in other words, is not between principles and values but between principles and preparedness.

This distinction matters because governance failures in AI deployment typically arise not from the absence of ethical intent but from the absence of institutional readiness. An institution may possess admirable governance aspirations while remaining unequipped to implement them when deployment actually occurs — a vulnerability that becomes acute in contexts characterised by resource constraints, evolving regulatory demands, and rapidly accelerating digital transformation agendas.

Existing frameworks address this problem only partially. Organisational readiness scholarship examines implementation capability but rarely focuses on governance legitimacy. Governance maturity models assess capability progression after structures exist rather than prerequisite conditions before deployment begins. Responsible AI frameworks specify normative expectations without attending systematically to whether institutions can meet them. The question that falls between these traditions — are institutions governance-ready before AI deployment occurs? — remains insufficiently explored.

This paper addresses that question through the concept of Ethical Readiness: the institutional governance condition in which accountability structures, governance ownership mechanisms, oversight capability, transparency preparedness, corrective governance capacity, governance adaptability, and governance proportionality are sufficiently developed to support responsible AI deployment. Rather than replacing existing responsible AI frameworks, Ethical Readiness functions as a complementary governance layer situated between ethical principles and operational deployment, focusing not on what institutions should do but on whether they are prepared to do it.

The paper contributes to the literature in four ways. First, it introduces Ethical Readiness as a distinct conceptual construct within responsible AI governance scholarship. Second, it differentiates this construct from adjacent concepts including governance maturity, organisational readiness, and responsible AI frameworks. Third, it operationalises the concept through a multidimensional governance architecture tailored to higher education. Fourth, it demonstrates practical relevance through comparative institutional evidence from contemporary university AI governance initiatives.

2. Literature Review

2.1 Responsible AI Governance

Responsible AI governance has evolved rapidly from a primarily philosophical concern into a multidisciplinary field encompassing public policy, organisational accountability, risk management, and technology regulation. The major international initiatives share several foundational assumptions: that institutions deploying AI must establish meaningful accountability mechanisms; that transparency and explainability are organisational as much as technical obligations; that governance structures must enable risk identification, mitigation, and corrective intervention; and that institutions remain responsible for the consequences of AI-supported decisions regardless of system sophistication.

Critical scholarship has usefully complicated this picture. Crawford (2021) and Mittelstadt et al. (2016) draw attention to structural power dynamics and algorithmic harm that normative principles may obscure. Selwyn (2019) raises questions about the implications of AI-driven systems for educational equity and professional autonomy. The EU Artificial Intelligence Act (European Parliament and of the Council, 2024) represents a significant regulatory development with direct implications for university governance across European systems; Veale and Borgesius (2021) provide an important analytical account of its governance implications.

What most of these frameworks share, however, is a focus on governance principles, objectives, or post-deployment outcomes rather than institutional preparedness prior to deployment. The question of whether institutions actually possess the governance capability to implement normative requirements — as opposed to merely committing to them — has received comparatively little systematic attention. This observation creates the analytical opening for Ethical Readiness as a governance-precondition construct.

2.2 Higher Education Governance and AI Adoption

Governance within universities differs from commercial governance in important structural respects. Universities operate within environments characterised by public accountability, academic autonomy, stakeholder diversity, and contested institutional legitimacy. Governance decisions affect not only operational outcomes but educational fairness, academic opportunity, student trust, and societal credibility. This environment, as Trow (1996) argued, demands governance mechanisms capable of preserving both accountability and academic values simultaneously.

The adoption of AI in higher education intensifies these demands. Universities now experiment with learning analytics, predictive student success models, AI-assisted assessment, intelligent tutoring systems, and administrative automation. Recent governance initiatives from institutions including the University of Edinburgh, Arizona State University, the University of Melbourne, and the National University of Singapore demonstrate growing recognition that effective AI governance requires institutional coordination rather than isolated technical controls. Nevertheless, approaches remain highly variable — some institutions emphasising policy development, others operational guidance, others awareness building — suggesting the absence of a consistent governance-preparedness framework.

2.3 Organisational Readiness and Governance Maturity

Weiner's (2009) theory of organisational readiness emphasises collective commitment and implementation capability as critical determinants of successful change; subsequent digital transformation research has similarly highlighted the importance of leadership support, coordination, and stakeholder engagement. These perspectives offer genuine insight into institutional preparedness but do not engage systematically with AI-specific governance obligations around accountability, oversight, and transparency. Institutional capacity frameworks — focused on operational capability, resource availability, and organisational resilience — share this limitation.

Governance maturity models such as COBIT provide mechanisms for assessing governance progression and benchmarking capability over time. Their limitation in the present context is temporal: they evaluate governance development after governance structures already exist rather than the governance conditions that must precede deployment. An institution may demonstrate mature governance in certain operational domains while remaining insufficiently prepared for AI-specific accountability and oversight challenges.

Ethical Readiness builds upon these traditions while maintaining a distinct analytical focus. The critical distinction is that readiness, in the present sense, is concerned with governance prerequisites — the conditions that must hold before deployment begins — rather than governance performance or capability progression after structures are established.

2.4 Research Gap

The reviewed literature demonstrates that existing frameworks address distinct and important governance concerns. Yet none directly focuses on institutional governance readiness as a prerequisite for responsible AI deployment. Responsible AI frameworks prioritise normative principles; organisational readiness models emphasise implementation capability; governance maturity approaches assess post-deployment progression; and higher education governance scholarship identifies institutional challenges without providing governance-precondition assessment tools. Table 1 summarises these distinctions.

Table 1. Governance Frameworks and Their Relationship to Pre-Deployment Readiness

Perspective	Primary Focus	Strength	Limitation
Responsible AI Governance	Ethical and governance principles	Normative guidance	Limited preparedness assessment

Organisational Readiness	Change implementation capability	Transformation support	Limited governance focus
Institutional Capacity	Organisational capability	Operational effectiveness	Limited AI governance emphasis
Governance Maturity	Governance optimisation	Capability assessment	Post-deployment orientation
Ethical Readiness	Governance preparedness before deployment	Governance-precondition evaluation	Requires empirical validation

3. Methodological Positioning

This study adopts a conceptual-theoretical research design aimed at developing and refining a governance framework for responsible AI deployment in higher education. Following Jaakkola's (2020) typology of conceptual article designs, the paper pursues a framework-building approach: its primary purpose is to introduce, define, and internally validate a new governance construct rather than test propositions against empirical data. This approach is consistent with the theory-building function identified by Gregor (2006), in which the construction and clarification of new explanatory frameworks constitutes a legitimate and distinct mode of scholarly contribution. Given the rapid evolution of AI governance practices and the limited availability of established governance-preparedness models for higher education, a conceptual-theoretical design is appropriate for the objectives of this study.

The Ethical Readiness framework was developed through integration of five scholarly domains: responsible AI governance literature provides foundational principles; higher education governance scholarship contributes perspectives on institutional legitimacy and accountability; organisational readiness research informs discussions of preparedness and institutional adaptation; governance maturity and institutional governance theories illuminate governance structures and capability development; and

institutional legitimacy theory — particularly the work of Suchman (1995), Meyer and Rowan (1977), and DiMaggio and Powell (1983) — provides the conceptual foundation for understanding why governance preparedness carries institutional significance beyond regulatory compliance.

Although the paper does not employ primary empirical data collection, it incorporates comparative institutional evidence derived from publicly documented AI governance initiatives implemented by universities across different educational and regulatory environments. These examples serve as analytical illustrations demonstrating how governance challenges identified in the literature are already being addressed through emerging institutional practices, thereby strengthening practical relevance while maintaining the conceptual orientation of the study.

Several important limitations follow from this design. The framework proposed here is a conceptual governance model rather than a validated measurement instrument. It does not claim universal applicability across all educational systems or regulatory environments, and the governance assumptions embedded within it draw primarily on Western and OECD-origin scholarship, a constraint that future empirical research should explicitly address. The eight governance dimensions are not asserted to be exhaustive, mutually exclusive, or universally applicable in identical form.

4. Conceptual Positioning of Ethical Readiness

Figure 1: Ethical Readiness as a Pre-Deployment Institutional Governance Layer

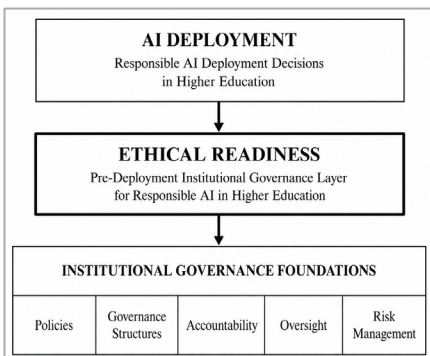


Figure 1 situates Ethical Readiness as a governance-precondition layer positioned between the external responsible AI governance ecosystem and the operational AI deployment environment. The positioning makes visible a governance transition zone that existing frameworks tend to overlook: not what institutions should do, but whether they are prepared to do it before AI systems become operationally embedded.

Ethical Readiness is defined as the institutional governance condition in which accountability structures, governance ownership mechanisms, oversight capability, transparency preparedness, corrective governance capacity, governance adaptability, and governance proportionality are sufficiently developed to support responsible AI deployment. This definition extends beyond policy compliance or ethical aspiration. An institution may publicly endorse responsible AI principles while remaining insufficiently prepared to implement them. Ethical Readiness is therefore concerned with institutional capability rather than institutional intention alone.

The concept rests on five foundational propositions: responsible AI governance requires institutional preparedness as a prior condition; governance capability is analytically distinct from governance intention; accountability mechanisms must be operationalised before deployment rather than retrofitted after problems emerge; governance intensity should vary proportionally with institutional context and risk exposure; and governance legitimacy depends upon both governance design and the readiness to execute that design.

Table 2 places Ethical Readiness in explicit comparative relation to existing governance models. The comparison reveals that Ethical Readiness does not displace these frameworks but occupies a governance-preparedness space that each leaves underdeveloped: governance maturity models focus on optimisation and progression; organisational readiness frameworks examine implementation capability; responsible AI frameworks establish normative expectations. Ethical Readiness differs by attending specifically to the governance conditions required before deployment occurs.

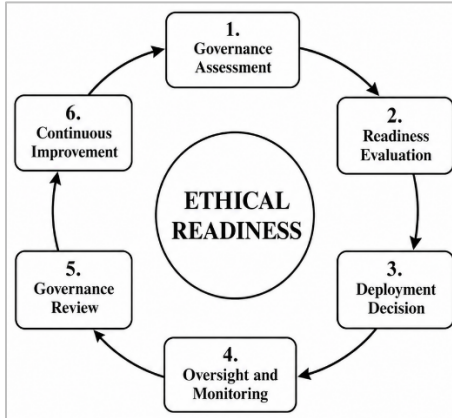
Table 2. Ethical Readiness Compared with Existing Governance Models

Governance Feature	Ethical Readiness	Governance Maturity	Org. Readiness	Responsible AI Frameworks
Pre-deployment focus	✓	✗	Partial	✗
Governance legitimacy	✓	Limited	Limited	Partial
Accountability mapping	✓	Partial	Partial	Partial
Governance proportionality	✓	✗	✗	Limited
Educational context sensitivity	✓	✗	✗	✗
Transparency readiness	✓	Partial	Limited	✓
Corrective governance capacity	✓	Limited	Limited	Partial
Oversight preparedness	✓	Partial	Partial	Partial

From this positioning, governance effectiveness depends upon three interdependent layers: governance principles specifying what institutions should do; governance readiness establishing whether institutions are prepared to do it; and governance execution determining how requirements are operationalised in practice. Ethical Readiness addresses the second of these layers — the layer most consistently absent from existing scholarship.

5. Ethical Readiness Dimensions

Figure 3: Ethical Readiness Dimensions Framework



The Ethical Readiness framework is operationalised through eight interdependent governance dimensions. These dimensions should be understood as interconnected governance capabilities rather than discrete controls: weakness in any single dimension reduces overall governance effectiveness even when other elements appear robust. The following discussion organises the dimensions into three analytical groupings, but this structure is analytical rather than hierarchical — all dimensions are operationally interdependent.

5.1 Governance Structure Dimensions

The most foundational governance requirements for responsible AI deployment concern the clarity of institutional rules, roles, and accountabilities. Without these structural prerequisites, more sophisticated governance mechanisms cannot function effectively.

Normative Clarity refers to the degree to which an institution has established clear expectations regarding acceptable and unacceptable AI-related practices. It functions as the common governance reference point from which decision-making, accountability, and institutional behaviour proceed. Its

absence is particularly consequential in higher education, where AI applications touch on academic assessment, data use, research integrity, and student rights: without shared normative expectations, inconsistency across faculties, departments, and administrative units becomes the default condition.

Role Clarity concerns the allocation of governance responsibilities, accountability ownership, oversight duties, and decision-making authority across institutional stakeholders. This dimension is especially challenging in universities because AI governance necessarily involves faculty members, administrators, technology teams, ethics committees, legal advisors, and institutional leadership simultaneously. When these roles overlap or remain unspecified, governance coordination degrades — not through bad intent but through structural ambiguity. Accountability Mapping extends this logic to traceability: it involves the explicit identification and documentation of governance ownership across the AI deployment lifecycle, from procurement and implementation through monitoring, review, escalation, and corrective intervention. Together, normative clarity, role clarity, and accountability mapping establish the governance architecture without which oversight and transparency obligations cannot be discharged.

5.2 Governance Control Dimensions

Where structural dimensions establish governance architecture, control dimensions enable active governance performance throughout the deployment lifecycle.

Oversight Capacity refers to the institutional ability to monitor, review, evaluate, and supervise AI-enabled activities. Oversight mechanisms provide governance visibility, enabling institutions to identify emerging concerns before they become operational failures. In higher education contexts — where AI is applied to assessment outcomes, student profiling, and academic progression — governance blind spots generated by inadequate oversight can carry direct educational consequences. The distinction between oversight as a formal commitment and oversight as a functioning capability is precisely the gap that Ethical Readiness targets.

Corrective Governance Capacity addresses what happens when oversight reveals problems. Many governance discussions focus on detection; Ethical Readiness insists equally on remediation. The capability to intervene — to suspend, review, investigate, remediate, and adapt when governance deficiencies are identified — is distinct from the capability to detect. Institutions may possess monitoring structures while lacking the authority, procedures, and coordination capacity to act on what monitoring reveals. Without corrective capacity, governance collapses into audit theatre.

5.3 Governance Adaptation Dimensions

The third grouping addresses the quality of governance judgement and the capacity of governance structures to evolve as technological, regulatory, and institutional environments change.

Contextual Judgement refers to the ability to calibrate governance intensity in proportion to institutional context, educational sensitivity, and risk exposure. Not all AI applications require identical governance interventions: assessment systems affecting academic progression demand substantially stronger accountability structures and transparency safeguards than scheduling automation. Uniform governance approaches either over-regulate low-risk deployments or, more dangerously, apply the same light-touch governance to high-stakes contexts as to routine administrative tools. Contextual judgement is what allows institutions to deploy governance resources where they create most value.

Transparency Readiness concerns institutional preparedness for governance disclosure, stakeholder communication, explainability, and decision visibility. Within universities, expectations of transparency extend from students seeking to understand how AI influences their academic assessment, to faculty members whose professional judgements may be mediated or supplemented by algorithmic systems, to regulators assessing institutional compliance with emerging AI legislation. Transparency readiness is not simply a commitment to disclosure — it requires the governance infrastructure through which disclosure becomes operationally possible.

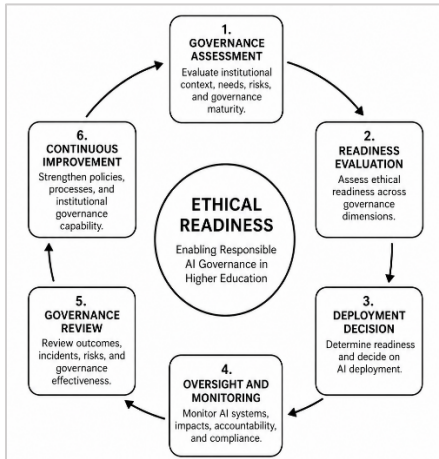
Governance Adaptability, finally, refers to the capacity of governance structures to evolve in response to technological change, shifting regulatory requirements, and emerging educational challenges. As generative AI continues to evolve, governance frameworks designed for earlier AI generations will require revision. Institutions that cannot adapt their governance — updating normative expectations, reallocating accountability, revising oversight arrangements — will progressively accumulate governance debt as deployed technologies outpace governance structures. Table 3 summarises all eight dimensions.

Table 3. Summary of Ethical Readiness Dimensions

Dimension	Governance Group	Primary Function	Risk if Weak
Normative Clarity	Structure	Establish governance expectations	Inconsistent institutional behaviour
Role Clarity	Structure	Allocate accountability	Responsibility diffusion
Accountability Mapping	Structure	Ensure lifecycle traceability	Accountability gaps
Oversight Capacity	Control	Enable governance visibility	Governance blind spots
Corrective Governance Capacity	Control	Enable intervention and remediation	Governance stagnation
Contextual Judgement	Adaptation	Support governance proportionality	Governance misalignment
Transparency Readiness	Adaptation	Strengthen trust and legitimacy	Stakeholder distrust
Governance Adaptability	Adaptation	Support continuous evolution	Governance obsolescence

6. The Ethical Readiness Governance Lifecycle

Figure 2: Ethical Readiness Governance Lifecycle Model



Ethical Readiness is not a threshold condition achieved at a single moment but a dynamic institutional capability that must be actively maintained across the full AI deployment lifecycle. Figure 2 presents a seven-stage model that clarifies how governance preparedness develops and operates over time.

The lifecycle divides analytically into two phases. Stages 1 through 3 constitute the pre-deployment readiness phase — the primary domain of the Ethical Readiness framework — addressing the governance conditions that must be established before AI systems become operationally embedded. Stages 4 through 7 represent governance-enabled operational activity, in which deployment proceeds within the accountability structures and oversight mechanisms established during the readiness phase. This distinction matters: Ethical Readiness does not govern deployment performance as such but establishes the architecture that makes responsible deployment possible.

The pre-deployment phase begins with Institutional Awareness: the recognition, at leadership and governance levels, that AI adoption introduces governance implications extending beyond technology implementation. Awareness is foundational because governance cannot address risks that have

not been recognised. This is followed by Governance Assessment, in which institutions examine existing governance structures against the eight Ethical Readiness dimensions, identifying capability strengths, weaknesses, and gaps before deployment occurs rather than after failures emerge. The third stage, Governance Structuring, transforms awareness and assessment into governance architecture: formalising accountability assignments, establishing oversight responsibilities, creating escalation pathways, and developing transparency expectations in ways that are traceable and enforceable.

The operational phase begins with Controlled AI Deployment: implementation under governance supervision, with continuous attention to accountability and oversight obligations. Stage 5, Monitoring and Oversight, sustains governance visibility throughout deployment — providing the information necessary for both correction and adaptation. Stage 6, Corrective Governance Intervention, addresses identified deficiencies through suspension, review, investigation, or remediation. The final stage, Institutional Learning, closes the loop: governance experience informs revisions to normative expectations, role allocations, oversight arrangements, and governance structures, strengthening preparedness for subsequent deployments. Table 4 summarises the lifecycle.

Table 4. Ethical Readiness Governance Lifecycle: Stages, Objectives, and Governance Functions

Stage	Phase	Governance Objective	Core Function
1. Institutional Awareness	Pre-deployment	Recognise governance obligations	Foundation for preparedness
2. Governance Assessment	Pre-deployment	Evaluate current readiness against dimensions	Gap identification
3. Governance Structuring	Pre-deployment	Establish governance architecture	Accountability and oversight design

4. Controlled Deployment	Operational	Deploy within governance boundaries	Governance-enabled implementation
5. Monitoring and Oversight	Operational	Maintain governance visibility	Continuous performance review
6. Corrective Intervention	Operational	Address identified deficiencies	Remediation and escalation
7. Institutional Learning	Operational	Adapt governance from experience	Continuous governance improvement

7. Institutional Governance Applications

The practical significance of Ethical Readiness is most visible when the framework is applied to specific AI deployment contexts. Different AI applications within universities generate substantially different governance demands, and the principle of governance proportionality — that governance intensity should be calibrated to educational sensitivity, accountability exposure, and risk severity — is one of the framework's most operationally important features.

7.1 AI-Assisted Assessment and Predictive Student Analytics

Assessment systems that incorporate AI to support grading, feedback generation, or academic progression decisions represent the most governance-sensitive deployment category within higher education. Their consequences extend directly to student outcomes and institutional credibility, making failures in accountability, transparency, or oversight particularly consequential. The governance requirements are correspondingly demanding: the highest levels of normative clarity, accountability mapping, transparency readiness, and oversight capacity are necessary to ensure that algorithmic influence on educational decisions remains explainable, traceable, and subject to meaningful human review. Institutions including the University of Melbourne and a number of UK universities have introduced AI assessment governance policies, disclosure requirements, and human review obligations that reflect this governance sensitivity in practice.

Predictive analytics systems — deployed to identify students at risk of underperformance, disengagement, or withdrawal — introduce related but distinct governance challenges. The potential benefits of proactive intervention are real; so are the risks of opaque profiling, algorithmic bias, and inadequate explainability. The governance question is not whether such systems should be used but whether institutions can implement them responsibly. Contextual judgement is particularly critical here, as are transparency readiness and accountability mapping: students whose academic trajectories are influenced by predictive models have legitimate interests in understanding the governance conditions under which those models operate.

7.2 Student Monitoring and Administrative Automation

Behavioural monitoring technologies — remote proctoring, engagement tracking, attendance monitoring, and activity surveillance — represent the AI deployment category most directly implicated in questions of educational trust and institutional legitimacy. The governance concerns are fundamental: privacy, autonomy, and the relationship between surveillance and educational environment. These applications require the highest levels of transparency readiness and corrective governance capacity, not only because of their intrinsic governance sensitivity but because the consequences of deploying them without adequate governance preparation extend to student trust and institutional reputation. The University of Edinburgh's generative AI guidance and comparable transparency frameworks from other institutions suggest growing awareness that student-facing AI applications require governance infrastructure before deployment rather than retrospective justification.

Administrative AI automation — chatbots, scheduling support, resource allocation assistance, service management — presents a genuinely different governance profile. Its educational consequences are generally lower, and governance requirements are correspondingly less intensive. Yet the principle of governance proportionality should not be read as governance permissiveness: role clarity and governance adaptability remain important even in lower-risk deployments, particularly given the tendency of

administrative automation to expand in scope over time. The distinction between these two application categories illustrates why uniform governance approaches are inadequate: the governance structures appropriate for proctoring systems would be disproportionate for scheduling tools and insufficient for assessment systems.

7.3 Generative AI Research Assistance

The rapid diffusion of generative AI tools into academic research workflows creates governance challenges that are both novel and institution-wide. Questions of authorship, attribution, intellectual ownership, and research integrity resist straightforward resolution, and normative expectations are still forming across the sector. What the Ethical Readiness framework highlights is that this uncertainty itself constitutes a governance risk: institutions that deploy generative AI research tools without first establishing normative clarity and accountability mapping are exposed to integrity challenges they will struggle to address retrospectively. Recent guidance from the University of Edinburgh and comparable international institutions has begun to establish disclosure obligations and attribution expectations, but governance adaptability will remain critical as both the capabilities of generative AI tools and the scholarly community's responses to them continue to evolve.

7.4 Governance Intensity Across Deployment Contexts

Table 5 consolidates the comparative governance analysis across the five application domains, illustrating how Ethical Readiness requirements should be calibrated to deployment context. The assessment is indicative rather than prescriptive: it is intended to support institutional governance reflection rather than replace contextual judgement.

Table 5. Governance Intensity and Ethical Readiness Requirements by Deployment Context

AI Application	Governance Sensitivity	Critical Dimensions	Ethical Readiness Requirement
AI-Assisted Assessment	Very High	Accountability Mapping, Transparency Readiness, Oversight Capacity	Very High
Student Monitoring / Proctoring	Very High	Transparency Readiness, Corrective Capacity, Contextual Judgement	Very High
Predictive Student Analytics	High	Contextual Judgement, Accountability Mapping, Transparency Readiness	High
Generative AI Research Assistance	Medium-High	Normative Clarity, Governance Adaptability, Accountability Mapping	High
Administrative AI Automation	Medium	Role Clarity, Governance Adaptability, Oversight Capacity	Moderate

Several patterns emerge from this comparative analysis that reinforce the framework's theoretical claims. First, governance requirements are genuinely heterogeneous across deployment types: the governance architecture appropriate for one application would be disproportionate or insufficient for another. Second, accountability mapping and transparency readiness appear as critical dimensions across multiple high-sensitivity contexts, suggesting they function as foundational governance capabilities rather than context-specific requirements. Third, the institutions whose governance initiatives are

referenced — Edinburgh, Melbourne, Arizona State, the National University of Singapore — have each developed governance mechanisms that implicitly address the Ethical Readiness dimensions most salient to their deployment priorities, offering preliminary institutional evidence that the framework reflects real governance dynamics rather than purely theoretical concerns.

8. Discussion

8.1 Governance Preparedness as a Prerequisite, not a Consequence

The central implication of this study is the reframing of governance capability as a prerequisite for responsible AI deployment rather than a consequence of operational experience. This distinction is less intuitive than it might appear. Much contemporary AI governance discourse — including many institutional policy documents — implicitly assumes that governance structures develop through iterative deployment: organisations deploy, encounter problems, and respond by strengthening governance. There is a genuine argument for this approach in domains where risks are modest and learning costs are manageable. In higher education, where AI applications can directly influence student outcomes, academic progression, and institutional legitimacy, the costs of retrospective governance development are substantially higher.

The Ethical Readiness framework operationalises this argument by specifying the governance conditions that must hold before deployment rather than merely documenting what governance should look like in theory. Institutions may possess policies, ethical guidelines, governance committees, and compliance obligations while simultaneously lacking the preparedness to implement those mechanisms effectively. Recognising this gap — and measuring it through the eight dimensions — is the framework's primary practical contribution.

8.2 Governance Proportionality and Institutional Context

A second contribution concerns the principle of governance proportionality. Universities are not governance monoliths: their AI deployments differ in

educational sensitivity, stakeholder exposure, regulatory relevance, and corrective consequence. A governance framework that imposes identical requirements across assessment systems and scheduling automation misallocates governance resources while potentially under-governing the applications that most warrant scrutiny. The comparative analysis in Section 7 demonstrates that governance requirements vary substantially and systematically across deployment contexts, and that institutions can benefit from frameworks that make this variation explicit and actionable.

This principle has practical implications for how governance committees are structured and resourced. Rather than applying uniform review processes to all AI deployments, institutions would benefit from governance triage mechanisms that allocate oversight intensity in proportion to deployment sensitivity. This is not a counsel for governance minimalism but for governance efficiency — ensuring that the highest-risk deployments receive governance attention commensurate with their stakes.

8.3 Theoretical Contributions

The study contributes to responsible AI governance scholarship through four analytical moves. First, it introduces governance preparedness as a distinct object of analysis, separable from governance principles, governance compliance, and governance maturity. This expansion of the analytical scope of responsible AI governance research opens space for empirical investigation into institutional readiness conditions that existing frameworks assume rather than examine.

Second, the framework integrates theoretical traditions that have developed largely in parallel — responsible AI governance, higher education governance, organisational readiness, governance maturity, and institutional legitimacy theory — into a unified governance-precondition architecture. This integration provides conceptual resources that no single tradition could supply: institutional legitimacy theory (Suchman, 1995; Meyer and Rowan, 1977) explains why governance preparedness carries significance beyond regulatory compliance; organisational readiness scholarship (Weiner, 2009)

informs the lifecycle model; accountability theory (Bovens, 2007) grounds the accountability mapping and corrective capacity dimensions.

Third, the framework operationalises governance preparedness through identifiable dimensions that can support institutional assessment and, importantly, future empirical investigation. Fourth, by situating the framework specifically within higher education, the study addresses governance concerns — educational accountability, academic integrity, student trust, institutional legitimacy — that sector-neutral frameworks necessarily abstract away.

8.4 Limitations, Tensions, and Critiques

Several limitations require candid acknowledgement. The framework remains conceptual and has not undergone systematic empirical validation. The comparative institutional evidence serves as analytical illustration rather than formal case study analysis. The governance assumptions embedded in the framework draw primarily on Western and OECD-origin scholarship, and their applicability across non-Western higher education systems, resource-constrained institutional contexts, and diverse regulatory environments requires explicit empirical investigation.

The eight dimensions are not claimed to be exhaustive or mutually exclusive, and the governance intensity assessments in Table 5 are indicative heuristics derived from governance logic rather than measured empirical findings. They should be treated as starting points for institutional reflection rather than conclusions. There is also a genuine risk that Ethical Readiness could be adopted performatively — as a compliance exercise or governance checklist — rather than as a substantive institutional capability. Addressing this risk requires empirical research into how institutions actually engage with governance-readiness frameworks in practice, rather than how they report engaging with them.

Three alternative interpretations deserve acknowledgement. Some scholars may argue that existing responsible AI frameworks already address institutional preparedness through their accountability and transparency provisions. This critique has some force: the distance between the present

framework and existing approaches is one of analytical emphasis rather than categorical difference. The response is that analytical emphasis matters — frameworks that foreground governance principles without attending to governance prerequisites produce different institutional behaviours than frameworks that treat preparedness as a first-order concern. A second critique might read Ethical Readiness as organisational readiness theory applied to AI governance, without substantive novelty. The response here is that the governance-specific dimensions — accountability mapping, corrective capacity, transparency readiness — represent adaptations that existing readiness frameworks do not supply. A third critique concerns governance burden: governance-intensive frameworks may slow innovation or create administrative complexity disproportionate to their benefits. The proportionality principle in the framework is a direct response to this concern, but the tension between governance rigour and institutional agility remains genuine and context-dependent.

9. Conclusion

Artificial intelligence is becoming structurally embedded within the operational, academic, and strategic functions of higher education institutions. The governance frameworks that have emerged in response — international principles, national regulations, institutional policies — represent genuine progress. They have clarified normative expectations, focused regulatory attention, and raised institutional awareness of the obligations associated with AI deployment.

What these frameworks have addressed less systematically is the question of whether institutions are actually prepared to govern before deployment occurs. This paper has argued that the gap between responsible AI commitments and responsible AI practice is often a gap of governance preparedness: not a failure of intention but a failure of readiness. Institutions may possess policies, committees, and ethical declarations while lacking the accountability structures, oversight mechanisms, corrective capacities, and adaptive governance capabilities that would allow those commitments to function in practice.

Ethical Readiness was introduced as a governance-precondition framework for addressing this gap within higher education. Defined as the institutional governance condition in which accountability structures, governance ownership mechanisms, oversight capability, transparency preparedness, corrective governance capacity, governance adaptability, and governance proportionality are sufficiently developed to support responsible AI deployment, the framework operationalises governance preparedness through eight interdependent dimensions and a seven-stage lifecycle model. Its applicability across assessment systems, predictive analytics, behavioural monitoring, generative AI tools, and administrative automation was examined through comparative institutional evidence, illustrating both the heterogeneity of governance requirements and the underlying governance capabilities that recur across high-sensitivity contexts.

Three directions for future research follow from this work. The most urgent priority is empirical validation through institutional case studies: qualitative research examining how governance structures, accountability mechanisms, and oversight capabilities develop in practice prior to AI deployment, with attention to variation across institutional types, national regulatory environments, and resource levels. Second, the development of a validated Ethical Readiness assessment instrument — capable of measuring institutional governance preparedness across the eight dimensions through survey methods, expert panel assessment, or document analysis — would operationalise the framework in ways that enable comparative institutional research and longitudinal tracking of governance development. Third, future research should examine whether institutions with higher pre-deployment governance preparedness demonstrate more accountable, transparent, and correctively responsive deployment behaviour — moving Ethical Readiness from a descriptive framework to a predictive governance model with direct policy implications.

Universities that deploy AI without adequate governance preparation carry a distinctive kind of institutional risk. It is not only the risk of algorithmic failure or regulatory non-compliance, but of deploying consequential technologies in environments where the governance conditions for responsible oversight do not yet exist. Ethical Readiness is proposed as a

framework for understanding, assessing, and strengthening those conditions. Ultimately, responsible AI deployment requires institutions that are not merely committed to responsible governance but prepared to execute it.

Use of Generative Artificial Intelligence

Generative AI tools were used during the preparation of this manuscript to provide language refinement, editorial assistance, brainstorming support, and structural formatting assistance. The author retained full intellectual responsibility for all aspects of the work, including conceptual development, literature review and synthesis, framework design, analytical argumentation, interpretation of evidence, and the conclusions drawn. All substantive scholarly content, theoretical claims, and research contributions are the author's own. The use of AI tools for language and formatting assistance does not affect the originality, integrity, or intellectual ownership of the research presented here. This disclosure is provided in accordance with academic publishing standards for transparency in AI-assisted scholarly work.

Declaration of Competing Interests

The author(s) declare no competing interests.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

10. Bibliography

- Armenakis, A. A., Harris, S. G., & Mossholder, K. W. (1993). Creating readiness for organizational change. *Human Relations*, 46(6), 681–703.
- Bovens, M. (2007). Analysing and assessing accountability: A conceptual framework. *European Law Journal*, 13(4), 447–468.
- Cath, C. (2018). Governing artificial intelligence: Ethical, legal and technical opportunities and challenges. *Philosophical Transactions of the Royal Society A*, 376(2133), 20180080.

- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Dafoe, A. (2018). *AI governance: A research agenda*. Future of Humanity Institute, University of Oxford.
- DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147–160.
- European Parliament and of the Council. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L series.
- Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1).
- Gregor, S. (2006). The nature of theory in information systems. *MIS Quarterly*, 30(3), 611–642.
- Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines*, 30(1), 99–120.
- Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial intelligence in education: Promises and implications for teaching and learning*. Centre for Curriculum Redesign.
- ISACA. (2018). *COBIT 2019 framework: Governance and management objectives*. ISACA.
- International Organization for Standardization. (2023). *ISO/IEC 42001: Artificial intelligence management systems — Requirements*. ISO.
- Jaakkola, E. (2020). Designing conceptual articles: Four approaches. *AMS Review*, 10(1), 18–26.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.
- Meyer, J. W., & Rowan, B. (1977). Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology*, 83(2), 340–363.
- Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1–21.
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)*. NIST.
- Organisation for Economic Co-operation and Development. (2019). *OECD principles on artificial intelligence*. OECD Publishing.

- Popenici, S. A. D., & Kerr, S. (2017). Exploring the impact of artificial intelligence on teaching and learning in higher education. *Research and Practice in Technology Enhanced Learning*, 12(1), 22.
- Power, M. (1997). *The audit society: Rituals of verification*. Oxford University Press.
- Scott, W. R. (2008). *Institutions and organizations: Ideas and interests* (3rd ed.). Sage.
- Selwyn, N. (2019). *Should robots replace teachers? AI and the future of education*. Polity Press.
- Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist*, 57(10), 1510–1529.
- Stensaker, B., Maassen, P., Borgan, M., Oftebro, M., & Karseth, B. (2019). Strengthening governance in higher education: The role of institutional leadership and accountability. *Higher Education Policy*, 32(2), 173–191.
- Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review*, 20(3), 571–610.
- Trow, M. (1996). Trust, markets and accountability in higher education: A comparative perspective. *Higher Education Policy*, 9(4), 309–324.
- UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. UNESCO.
- Veale, M., & Borgesius, F. Z. (2021). Demystifying the draft EU Artificial Intelligence Act. *Computer Law Review International*, 22(4), 9112.
- Weiner, B. J. (2009). A theory of organizational readiness for change. *Implementation Science*, 4, Article 67.
- Whittlestone, J., Nyrup, R., Alexandrova, A., & Cave, S. (2019). The role and limits of principles in AI ethics: Towards a focus on tensions. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (pp. 195–200). ACM.
- Williamson, B. (2019). Policy networks, performance metrics and platform markets: Charting the expanding data infrastructure of higher education. *British Journal of Educational Technology*, 50(6), 2794–2809.
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education: Where are the educators? *International Journal of Educational Technology in Higher Education*, 16(1), 39.

11. Short biography

Mohamed Hussien Abdelaziz is an AI governance and institutional transformation specialist working at the intersection of artificial intelligence, governance, ethics, and higher education futures. His work focuses on responsible AI deployment, institutional accountability, and the governance of digitally transforming organisations in regulated and emerging contexts. He holds a Ph.D. in AI/NLP and is particularly interested in how AI reshapes epistemic environments, legitimacy, and decision-making across complex institutional settings.

Email: mh.abdelaziz@ejada.com