

Article

Not peer-reviewed version

From Innovation Signaling to Governance Readiness: Generative Artificial Intelligence Integration, Policies, Opportunities, and Challenges in Two Rwandan Universities

Fabrice Dusengumuremyi^{*}, Henry John Chukwudi, Sylvie Ndahimana, Claver Ndahayo, [Sixbert Sangwa](#)

Posted Date: 23 April 2026

doi: 10.20944/preprints202604.1703.v1

Keywords: generative artificial intelligence; higher education; Rwanda; academic integrity; AI governance; institutional readiness; document analysis; comparative case study



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

From Innovation Signaling to Governance Readiness: Generative Artificial Intelligence Integration, Policies, Opportunities, and Challenges in Two Rwandan Universities

Fabrice Dusengumuremyi ^{1,*}, Henry John Chukwudi ², Sylvie Ndahimana ³, Claver Ndahayo ⁴ and Sixbert Sangwa ³

¹ Department of Skills Immersion, African Leadership University, Kigali, Rwanda

² Library Services, African Leadership University, Kigali, Rwanda

³ Department of International Business and Trade, African Leadership University, Kigali, Rwanda

⁴ Office of the Deputy Vice-Chancellor, Adventist University of Central Africa, Kigali, Rwanda

* Correspondence: fdusengumuremyi@alueducation.com

Abstract

Background: Generative artificial intelligence is entering higher education faster than many universities have been able to govern it, particularly in African contexts where policy ambition, institutional capacity, digital infrastructure, and pedagogical practice do not always advance at the same pace. **Purpose:** This study examines how generative artificial intelligence integration is publicly documented, governed, and framed at two universities in Rwanda: the African Leadership University and the Adventist University of Central Africa. **Design:** Guided by an integrated framework combining institutional readiness, Diffusion of Innovations, and a rights-based governance lens, the study adopts an interpretivist comparative multiple-case design based on document analysis and secondary analysis. The corpus comprises publicly retrievable institutional webpages, policy documents, academic regulations, handbooks, e-learning materials, research manuals, national policy texts, and recent peer-reviewed scholarship published or available between 2021 and April 2026. **Findings:** Public evidence indicates visible AI engagement at both universities, but in materially different forms. ALU appears more innovation-signalling, foregrounding AI research, student bootcamps, and academic-support programming. AUCA appears more governance-dense, with stronger public visibility of academic regulations, academic-integrity language, ICT and online-learning policies, plagiarism infrastructure, and AI-and-big-data institutional positioning. However, neither institution publicly presents a fully specified generative AI acceptable-use regime aligned with Rwanda's evolving national and sectoral AI governance expectations. The findings therefore suggest that visible experimentation is advancing faster than visible rule specificity. **Originality/value:** The study contributes rare comparative African evidence on university AI governance and introduces a useful analytical distinction between innovation signalling and governance readiness. **Practical implications:** The central challenge is no longer whether universities will adopt AI, but whether they can align policy clarity, academic-integrity architecture, digital capacity, and educational purpose in institutionally credible ways. The study also identifies concrete priorities for later primary research on implementation, stakeholder interpretation, and assessment design.

Keywords: generative artificial intelligence; higher education; Rwanda; academic integrity; AI governance; institutional readiness; document analysis; comparative case study

1. Introduction

1.1. Background and Context

Generative artificial intelligence has rapidly altered the conditions under which university teaching, learning, assessment, and research now occur. UNESCO's global guidance treats generative AI as both an educational opportunity and a governance problem, emphasizing that human agency, transparency, equity, and policy capacity must accompany technological uptake. Recent higher education scholarship likewise shows that large language models and related tools are no longer peripheral; they are reshaping academic writing, feedback, assessment design, and the meaning of student authorship itself (Kasneci et al., 2023; Miao & Holmes, 2023).

1.2. Problem Statement

The most important scholarly problem is no longer whether generative AI will be used in universities, but whether institutions can govern its use with sufficient clarity, legitimacy, and pedagogical seriousness. Comparative studies of university policies show that many institutions initially responded to generative AI through rushed or fragmented guidance, often focused narrowly on cheating, plagiarism, or originality, while more recent policy language is moving toward managed integration, disclosure expectations, and broader stakeholder responsibilities (An et al., 2025; Jin et al., 2025; Luo, 2024; Moorhouse et al., 2023). Yet this literature remains dominated by Global North institutions, leaving African universities underdocumented in a domain where context clearly matters (Sangwa & Mutabazi, 2025).

1.3. Research Objectives and Questions

This article aims to examine how generative AI integration is publicly documented, governed, and framed at the African Leadership University and the Adventist University of Central Africa. Three objectives organize the study: to map documented AI-related platforms, uses, and institutional references; to compare policy clarity, governance mechanisms, and readiness conditions; and to identify opportunities, risks, ambiguities, and evidence gaps that should inform later primary-data research.

The study asks three questions: **RQ1**. What generative AI platforms, activities, or institutional references are publicly documented by ALU and AUCA? **RQ2**. How do the two universities compare in policy architecture, academic-integrity framing, infrastructure signals, and capacity-building provisions? **RQ3**. What opportunities, risks, and governance gaps emerge from the documentary record, and what remains unknowable without later fieldwork?

1.4. Significance, Scope, and Delimitations

The study is significant for three reasons. Empirically, it contributes scarce African evidence on university generative AI governance. Conceptually, it treats AI integration not as a purely technical matter but as a question of institutional readiness, ethical legitimacy, and pedagogical authority. Methodologically, it demonstrates that secondary-data case study research can produce a defensible first-stage article when claims are calibrated to documentary limits. The analysis is delimited to public and retrievable materials from 2021 to April 2026, with strongest emphasis on 2022 to 2026 because this period captures the post-ChatGPT acceleration of policy and guidance responses. The findings do not claim to measure prevalence of student use, staff attitudes, or classroom implementation fidelity (Bowen, 2009). This article contributes in three ways. First, it provides one of the few comparative analyses of publicly documented generative AI governance in Rwandan higher education under a shared national policy environment. Second, it introduces a useful analytical distinction between innovation signaling and governance readiness, showing that visible AI activity and visible AI rule architecture do not necessarily mature at the same pace. Third, it demonstrates the value of public-documentary analysis as a first-stage method for mapping institution-level AI

governance in underdocumented contexts before later primary research tests how these visible arrangements are interpreted, enacted, or resisted in practice.

2. Literature and Theoretical Framework

2.1. Conceptualising Generative AI in Higher Education

Generative AI in higher education is best understood as a socio-technical and institutional phenomenon rather than merely a new classroom tool. Recent reviews emphasize that these systems alter the distribution of cognitive labor, making it easier to externalize drafting, translation, summarization, and coding while intensifying disputes over authorship, accountability, and evidence of learning (Dwivedi et al., 2023; Kasneci et al., 2023; Lodge et al., 2023). The educational challenge is therefore not simply tool adoption, but the redesign of assessment, integrity norms, and pedagogical expectations under conditions of ubiquitous machine assistance.

2.2. Global Scholarship on Governance, Adoption, and Risk

The strongest recent policy literature identifies a transition from reactive prohibition to differentiated integration. Chan (2023) proposes a three-dimensional university AI policy framework spanning pedagogical, governance, and operational issues. An et al. (2025) show that institutional guidance is often distributed across multiple audiences and functions rather than consolidated in one coherent policy text. Jin et al. (2025), studying 40 universities across six global regions, find that universities increasingly frame generative AI in terms of academic integrity, teaching enhancement, and equity, but still require clearer roles, responsibilities, and communication strategies. Luo (2024) further argues that many university policies still misdiagnose the deeper problem by treating generative AI too narrowly as a threat to originality rather than a trigger for rethinking assessment. Belkina et al. (2025) add that implementation research remains less developed than commentary and policy rhetoric, which is precisely why institution-specific documentary analysis remains valuable (Bowen, 2009; Johnston, 2014). Table 1 synthesizes representative primary and review studies that are most relevant to the present article's comparative and governance-centered focus.

Table 1. Representative Recent Scholarship on Generative AI Governance in Higher Education and Its Relevance to the Present Study.

Study	Context and sample	Method	Main contribution	Relevance to this article
Chan (2023)	Hong Kong universities; students, teachers, staff	Mixed methods	Proposes pedagogical, governance, and operational AI policy framework	Supports multi-dimensional policy analysis
Moorhouse et al. (2023)	Top-ranked universities globally	Policy/document analysis	Shows rapid but uneven guideline development for assessment and GAI	Benchmarks policy variability
An et al. (2025)	Top 50 U.S. universities	Mixed methods and guideline analysis	Demonstrates differentiated guidance for faculty, students, research, and administration	Helps interpret distributed governance

Jin et al. (2025)	40 universities across six global regions	Comparative policy analysis using Diffusion of Innovations	Shows proactive but uneven global institutional adoption strategies	Anchors comparative theoretical lens
Luo (2024)	20 world-leading universities	Critical policy review	Argues policies over-focus on originality and under-address assessment redesign	Sharpens the article's critical discussion
Sebihi et al. (2025)	University of Rwanda	Survey-informed policy analysis	Identifies readiness potential alongside regulatory and infrastructure gaps	Establishes Rwandan higher education relevance
Belkina et al. (2025)	International case studies	Systematic review	Notes the field's move from commentary to implementation research	Justifies document-based first-stage analysis

The table is an author synthesis of recent peer-reviewed scholarship and is intended to show how the manuscript's contribution differs from prior work. Source: Authors' synthesis based on Chan (2023), Moorhouse et al. (2023), An et al. (2025), Jin et al. (2025), Luo (2024), Sebihi et al. (2025), and Belkina et al. (2025).

2.3. GAI in African and Global South Higher Education

The African literature remains much thinner than the literature from North America, Europe, Australia, and East Asia. Rwandan evidence from the University of Rwanda highlights optimism about AI's potential while underscoring infrastructure constraints, regulatory needs, and ethical concerns. Comparative work involving Rwanda, Lesotho, and Nigeria similarly argues that training, planning, and smart-technology provision remain central barriers to fuller adoption (Theodorio et al., 2024). An emerging African case study from the University of KwaZulu-Natal is significant because it shows that explicit institutional GenAI governance is possible on the continent, but still exceptional rather than ordinary (University of KwaZulu-Natal, 2025). This underdocumentation matters because governance assumptions imported from better-resourced universities do not automatically translate into African institutional contexts (Sebihi et al., 2025).

2.4. Rwanda's Policy and Higher Education Context

Rwanda's policy environment is unusually salient for this study. The National AI Policy seeks to position Rwanda as a center for responsible and inclusive AI, explicitly prioritizing 21st-century skills and AI literacy, reliable infrastructure and compute capacity, a robust data strategy, and practical ethical guidelines (MINICT, 2023). Crucially for higher education, the policy calls for world-class AI university education and applied research, research fellowships, graduate training, and long-term public-sector support to universities. By late 2025, the Rwandan Higher Education Council had gone further by issuing guidelines that requires higher learning institutions to align virtual-learning and AI practices with ethical integration, academic integrity, institutional capacity, and even AI-use declaration mechanisms for students, lecturers, and researchers (HEC, 2025). Rwanda therefore provides an analytically rich setting in which national ambition is explicit while institution-level implementation remains uneven and only partly visible (MINICT, 2023).

2.5. Institutional Readiness, Technology Acceptance, and Diffusion of Innovations

This article uses an integrated theoretical framework built around institutional readiness, Diffusion of Innovations, and a rights-based governance lens. Institutional readiness is treated as the primary explanatory construct and is operationalized through governance clarity, digital infrastructure, leadership attention, staff and student support, training, and ethical safeguards. This reading is consistent with OECD's account of effective digital education ecosystems and with UNESCO's recent frameworks on AI governance, teacher competency, and the protection of learners' rights in AI-mediated educational settings. Diffusion of Innovations adds an organizational and communication lens, clarifying why institutions differ in compatibility judgments, observability, trialability, and the public spread of new practices. Together, these lenses allow the study to distinguish between visible innovation signaling and visible governance readiness under conditions of rapid technological change.

The Technology Acceptance Model is not used as a primary framework in this article because its central constructs, especially perceived usefulness and perceived ease of use, are more appropriately tested through interviews, surveys, or other primary data on user perceptions. In the present document-based design, a rights-based governance lens offers a stronger fit because the evidence consists of publicly retrievable policy texts, platform signals, and institutional communications rather than user-level attitudinal data.

2.6. Conceptual Gap

Existing literature is now rich in commentary, policy scans, and perception studies, but it still lacks sufficient African comparative case studies that combine institutional documents, national policy context, and readiness theory in a disciplined way. The present article addresses that gap by comparing two Rwandan universities under one national policy environment while carefully avoiding claims that documentary data cannot sustain (Belkina et al., 2025).

3. Methodology

3.1. Research Paradigm and Design

The study adopts an interpretivist qualitative paradigm and a comparative multiple-case study design based on document analysis and secondary analysis. An interpretivist stance is appropriate because the study is not estimating prevalence or causal effect; it is interpreting how institutions publicly signify AI, formalize governance, and render academic expectations legible through documents. Following Bowen (2009), documents are treated not as weak substitutes for "real" data but as evidence of formal institutional position, public communication, and governance architecture. The design is appropriate because the research questions concern visible policy articulation, formal readiness signals, and publicly framed opportunities and risks.

3.2. Case Selection and Data Sources

ALU and AUCA were selected because they are both universities in Rwanda but differ in mission identity, public-facing digital ecosystems, and visible AI-related activity. The documentary corpus comprised public institutional webpages, policies, regulations, handbooks, e-learning materials, research manuals, national Rwandan AI and higher education policy texts, and recent peer-reviewed scholarship. For ALU, the most analytically important public documents were the ALU policies page, the AI in Africa research-project pages, the AI bootcamp page, the mostly digital library portal, and the library's academic-integrity workshop announcement (ALU, n.d.-a, n.d.-b, n.d.-c, n.d.-d). For AUCA, the core materials were the academic regulations, student handbook, ICT policy, online teaching and learning policy, research manual, research center page, homepage AI outreach material, and the Moodle-based e-learning platform (AUCA, 2021a, 2021b, 2022, 2023, n.d.-a, n.d.-b,

n.d.-c, n.d.-d). National benchmark documents were Rwanda's National AI Policy and the Higher Education Council's 2025 guidelines on virtual learning and AI (MINICT, 2023; HEC, 2025)..

3.3. Inclusion, Exclusion, and Analytical Procedure

Documents were included when they were institutionally issued or hosted, retrievable, and relevant to teaching, learning, assessment, digital infrastructure, academic integrity, AI, or related governance concerns. Duplicates and purely promotional materials without analytic substance were excluded, although public-facing communications were retained when they documented identifiable AI activity or institutional priorities. Analysis proceeded in four stages: document appraisal; structured extraction into a case-comparison matrix; directed qualitative content analysis using readiness, acceptance, and diffusion codes; and cross-case synthesis. Hsieh and Shannon's directed content analysis was used because the goal was not open-ended thematic discovery alone, but theory-informed comparison (Hsieh & Shannon, 2005).

Directed coding used six domains: public AI visibility, public academic-integrity specificity, public digital-governance density, AI-specific acceptable-use specificity, capacity-building signals, and national-policy alignment visibility. A document contributed to a code only when it contained directly retrievable language, policy provisions, or program descriptions relevant to that domain. Promotional references were analytically distinguished from binding governance texts and were not treated as equivalent forms of evidence. Cross-case comparison was then conducted at the level of coded documentary patterns rather than assumed institutional practice. The documentary workflow below summarizes the research sequence used.

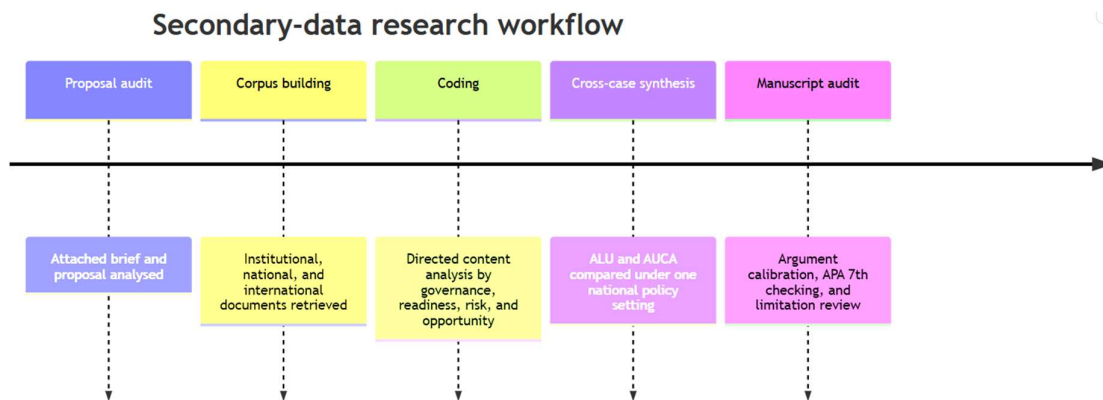


Figure 1. Documentary Analysis Workflow for the Comparative Multiple-Case Study This timeline is an author-generated representation of the method followed in the present study and is grounded in the documentary corpus assembled for the cases.

3.4. Trustworthiness, Ethics, and Design Limits

Trustworthiness was strengthened through source triangulation across institutional, national, and international documents; a cross-case coding matrix; and explicit separation of direct documentary evidence from analytic inference. The study involved no human participants and relied on public documents, so ethical risk was minimal. The main limitation is intrinsic to the design: public documents cannot reveal the full gap between formal policy and lived practice. Nor can public absence prove institutional absence. What the study can show, however, is the visible governance architecture that reviewers, staff, students, and external stakeholders can actually find (Bowen, 2009).

4. Findings and Discussion

4.1. Overview of the Documentary Corpus

The institutional corpus yielded materially different public documentation ecologies. ALU's retrievable corpus was smaller and more outward-facing, centered on AI ambition, academic support, and broad institutional ethics. AUCA's corpus was larger and more rule-heavy, including academic, ICT, online-learning, integrity, research, and LMS materials. This difference is substantively important because a university's public document ecology is itself part of how governance is communicated and experienced. Table 2 summarizes the core documentary corpus used in the cross-case analysis.

Table 2. Core Documentary Corpus Used in the Cross-Case Analysis.

Corpus cluster	Institution or source	Document type	Primary analytical use
ALU policies page	ALU	Institutional webpage	Public governance visibility
AI in Africa project and forum pages	ALU	Research and policy webpages	AI ambition and governance discourse
Synapse Africa AI Bootcamp page	ALU	Student-program webpage	AI capacity-building signal
Library home and upcoming events pages	ALU	Library support webpages	Digital support and integrity programming
Academic regulations, student handbook, ICT policy	AUCA	Formal policy PDFs	Governance, conduct, and digital-use architecture
Online teaching and learning policy, Moodle portal	AUCA	Online-learning policy and LMS	Readiness and digital-learning support
Research manual, research center, AI outreach pages	AUCA	Research/governance webpages and PDF	Research culture, AI innovation, publication expectations
National AI policy and HEC AI guidelines	Rwanda	Official national policy documents	Macro-policy benchmark and sector expectations

This table presents the institutional and national documents included in the study, identifying each corpus cluster, institution or source, document type, and primary analytical use in the comparison of ALU and AUCA.

4.2. Documented GAI Platforms, Uses, and Institutional References

ALU's public-facing AI footprint is substantial but mostly developmental rather than regulatory. Its AI in Africa Research Project states that ALU launched a cross-country initiative to inform policymakers, businesses, and civil society on responsible AI development and deployment in Sub-Saharan Africa, explicitly covering education among its focal sectors (ALU, n.d.-b). ALU also advertises a Synapse Africa AI Bootcamp in partnership with Colorado State University and the Global Livingston Institute, designed to give students hands-on exposure to artificial intelligence (ALU, n.d.-c). These documents show that AI is not marginal at ALU; it is publicly normalized as a research, learning, and support topic.

AUCA's public-facing AI footprint is equally visible, but it is embedded more directly in institutional structures. AUCA's homepage advertises an "AI & Big Data training outreach program," linking AI activity to Rwanda's wider national digital-skills agenda (AUCA, n.d.-a; One Million

Rwandan Coders, n.d.). AUCA's research center page identifies AI and big data as one of its strategic research areas and states that its innovation center works on AI and big-data applications and training linked to healthcare, education, and agriculture (AUCA, n.d.-d). The AUCA e-learning environment exposes operational digital infrastructure through policy links, a plagiarism link to Turnitin, an active course titled "plagiarism check," and a postgraduate Big Data Analytics category (AUCA, n.d.-b). Together, these materials depict AUCA as publicly positioning AI not merely as an external trend but as part of its innovation, postgraduate, and platform ecosystems.

The difference between the cases is therefore not visibility versus invisibility, but the form of visibility. ALU's public AI presence is outward-facing, partnership-rich, and oriented toward regional dialogue and student exposure. AUCA's is more institution-internalized, linking AI to formal digital systems, postgraduate training, and research-center identity. This distinction supports a nuanced comparison: AI visibility alone does not tell us whether governance clarity is equally mature. Figures 2 and 3 visualize the frequency with which documents in each case contributed to the predefined coding domains described in Section 3.3; because categories were non-mutually exclusive, a single document could contribute to more than one coded domain.

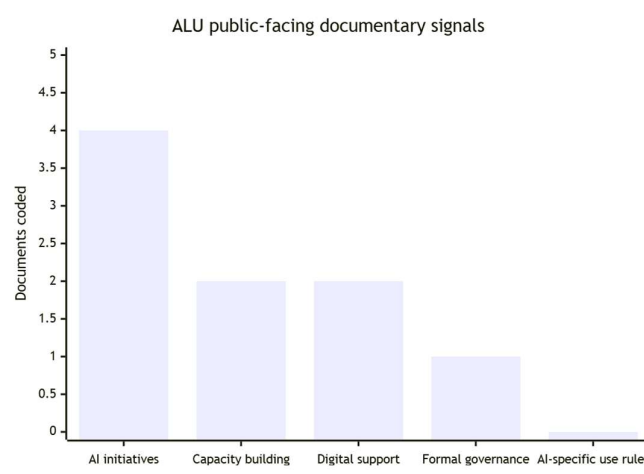


Figure 2. Distribution of Coding-Domain Hits in the ALU Documentary Corpus. This chart summarizes non-mutually-exclusive category hits in the ALU document set assembled for this study. It indicates strong public AI signaling and weaker public rule specificity.

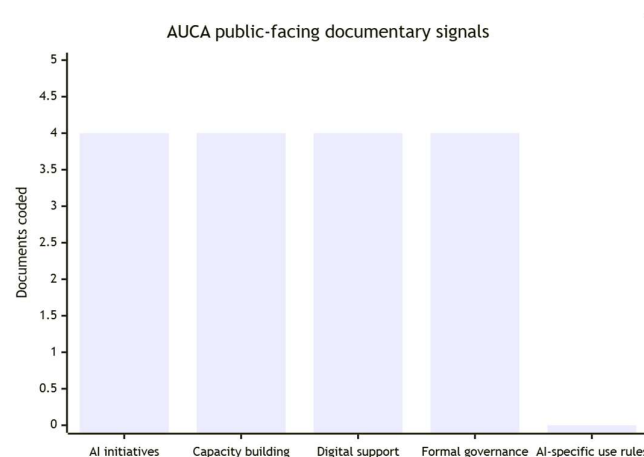


Figure 3. Distribution of Coding-Domain Hits in the AUCA Documentary Corpus. This chart summarizes non-mutually-exclusive category hits in the AUCA document set assembled for this study. It shows a denser public governance and readiness environment than ALU's, but also the same absence of a clearly retrievable institution-specific generative AI acceptable-use regime.

4.3. Publicly Visible Policy Architecture at ALU and AUCA

The most consequential difference lies in policy architecture. ALU's official policies page publicly lists a code of ethics, safeguarding policy, conflicts-of-interest policy, anti-bribery and anti-fraud policy, and whistleblowing policy (ALU, n.d.-a). These are meaningful governance documents, but the public page does not visibly provide AI-specific or academic-integrity-specific rules for classroom use of generative AI. ALU's library and program pages show academic support and AI engagement, but public-facing guidance on disclosure, permitted assistance, or course-level acceptable use was not retrievable in the corpus assembled for this article. That does not establish institutional absence, but it does establish public under-visibility.

AUCA's public policy architecture is thicker. Its academic regulations formalize teaching, learning, assessment, and student responsibilities (AUCA 2023, 2021b, n.d.-c, 2021a, n.d.-b). Its ICT policy establishes university-wide governance over ICT, software, cybersecurity, service support, and capacity building. Its online teaching and learning policy emerged from explicit efforts to clarify digital teaching arrangements. Most importantly for academic integrity, AUCA's student handbook spells out academic dishonesty, defines plagiarism to include electronic sources, and provides sanctions ranging from zero grades to suspension or dismissal in serious cases. AUCA's Moodle environment also links users directly to e-learning and blended-learning policies and to plagiarism infrastructure. Taken together, AUCA makes more of its academic and digital governance publicly legible.

Yet AUCA's greater governance density does not mean full generative AI policy maturity. The rules that are public are mostly legacy academic-integrity and digital-governance texts rather than explicit generative AI acceptable-use regimes. This matters because the Higher Education Council's 2025 guidelines now expect institutions to specify responsible AI use, map permission levels, and use AI declaration mechanisms (HEC, 2025). By that benchmark, both universities remain publicly incomplete, though AUCA starts from a more substantial rule base

Table 3. Cross-Case Comparison of Publicly Visible Generative AI Governance Readiness at ALU and AUCA.

Analytical domain	ALU	AUCA	Comparative interpretation
Public AI visibility	High through research project, forum, bootcamp, workshop	High through outreach, research center, LMS, big data identity	Both visibly engaged with AI
Public academic-integrity specificity	Limited in retrievable corpus	Clear plagiarism and dishonesty rules in handbook	AUCA stronger on visible integrity architecture
Public digital-governance density	Thin on retrievable academic digital rules	Stronger through ICT, online-learning, and LMS policy links	AUCA denser rule environment
AI-specific acceptable-use rules	Not publicly retrievable	Not publicly retrievable	Shared governance gap
Capacity-building signals	AI bootcamp; integrity workshop; digital library	AI outreach; research center; Moodle; big data postgraduate pathway	Both show readiness signals, differently distributed

National-policy alignment visibility	Implicit and innovation-oriented	More institutionally internalized but still incomplete	Neither yet visibly aligned in full to HEC 2025 model
Overall public-facing readiness profile	Innovation-forward, governance-thin	Governance-denser, still AI-rule incomplete	Different strengths, same missing endpoint

This table compares the two universities across key analytical domains, including public AI visibility, academic-integrity specificity, digital-governance density, AI-specific acceptable-use rules, capacity-building signals, national-policy alignment visibility, and overall public-facing readiness profile. **Source:** Authors' cross-case synthesis of the documentary corpus in Table 2, interpreted against Higher Education Council (2025).

4.4. Publicly Visible Readiness Signals

ALU's readiness signals are strongest in infrastructure and intellectual orientation. The ALU library describes itself as "primarily digital" for its Rwanda students and reports access to more than 263,300 journals (ALU, n.d.-d). ALU's software engineering page also frames the program around ethical digital transformation. These are not trivial claims. They suggest an environment in which digital resources, technological identity, and ethical language are already part of institutional self-description. In readiness terms, ALU appears to possess significant symbolic and infrastructural capacity, even if public governance specificity lags behind public AI ambition.

AUCA's readiness signals are more operational. The ICT policy explicitly covers governance, data communications, cybersecurity, software acquisition, service management, skills capacity building, and data warehousing. The e-learning platform provides policy links, course infrastructure, plagiarism-checking functionality, and postgraduate analytics-related offerings; because the visible user and course totals are dynamic webpage indicators, they are treated here as dated platform signals rather than stable institutional metrics. The research manual further signals an institutional research culture, recognizing documentary and historical studies as legitimate research forms and explicitly stating that AUCA uses APA style. These materials point toward a more procedurally articulated readiness environment, even if not yet translated into explicit public generative AI governance.

Readiness, however, is not only technical. It is also institutional self-knowledge. UNESCO and OECD both warn that digital transformation fails when institutions treat technology acquisition as sufficient and neglect human capacity, governance, and trust (UNESCO, 2023; OECD, 2023). On this criterion, both cases are only partially ready. ALU's public ecosystem suggests high perceived usefulness and experimentation, consistent with Technology Acceptance Model and Diffusion of Innovations expectations, but public policy visibility is insufficiently commensurate with that experimentation. AUCA's environment suggests stronger administrative and platform readiness, but its public rules still do not clearly specify how generative AI may be used across assignment types or how disclosure should operate.

4.5. Governance Gaps and Areas of Ambiguity

Three governance gaps stand out. The first is the absence of publicly visible institution-specific generative AI acceptable-use matrices at either university. The Higher Education Council's 2025 guidelines are notable because they explicitly reference red, amber, and green permission logics for AI use and include an AI-use declaration form for students, lecturers, and researchers. No comparably explicit public guidance was retrievable for ALU or AUCA (HEC, 2025).

The second gap concerns the movement from generic integrity language to AI-era integrity language. AUCA does define plagiarism and academic dishonesty, including the misuse of electronic sources (AUCA, 2021b), which is a stronger base than what was publicly retrievable for ALU. But neither case publicly clarified, in the corpus used here, how disclosure, verification, citation of AI

assistance, or differential permission across assessment types should be handled in the generative AI era. Scholarship suggests that this is now the crucial policy frontier, because the key issue is not simply copied text but the regulation of machine-assisted intellectual labor (Luo, 2024).

The third gap is public transparency about faculty and student support. ALU and AUCA both show capacity-building signals, yet public documentation does not clearly establish a systematic, university-wide program for staff development, curriculum redesign, or AI literacy certification comparable to the ambition expressed in Rwanda's national AI policy. This matters because Rwanda's macro-policy now explicitly links AI leadership to university capacity building, graduate-level training, and long-term institutional strengthening (MINICT, 2023; HEC, 2025). Public-facing under-specification at university level therefore creates a visible implementation gap between national aspiration and campus rule architecture.

4.6. Opportunities for Teaching, Learning, and Institutional Innovation

The documentary record also reveals real opportunities. For ALU, AI is publicly tied to research leadership, student exposure, ethical discourse, and technological identity. That combination could support rapid movement into mission-aligned AI literacy, assessment redesign, and regionally influential governance development. For AUCA, the combination of formal digital governance, LMS infrastructure, AI outreach, and a big-data academic profile provides a stronger procedural foundation for institutionalizing AI guidance within teaching and learning systems. In both cases, the public surface shows that generative AI can be framed not only as a threat to integrity but as a catalyst for better digital pedagogy, research support, and graduate capability formation. This aligns with broader literature that sees value in generative AI when it is pedagogically bounded, transparent, and critically supervised.

4.7. Risks, Ethical Tensions, and Practical Constraints

The risks are equally clear. UNESCO warns that generative AI can undermine human agency, equity, privacy, and trustworthy knowledge if institutions fail to regulate use and build capacity. OECD similarly stresses that effective digital ecosystems depend on trustworthy governance, not just access (UNESCO, 2023; OECD, 2023). In the present cases, the primary risk is misalignment: public AI promotion may normalize use faster than formal rules can clarify responsibility. Where that occurs, universities risk creating a moral economy of ambiguity in which students and staff either overuse AI, underreport it, or operate under inconsistent assumptions across courses and units. This is not merely an administrative inconvenience. It is a philosophical problem about truthfulness, epistemic formation, and the integrity of educational judgment.

The deeper issue is not only whether students use AI, but whether assessment still warrants the inferences that universities claim to draw from submitted work. Once generative systems participate in drafting, summarization, translation, or ideation, questions of authorship become inseparable from questions of epistemic responsibility, disclosure, and the validity of assessment as a judgment of student learning. A governance framework that treats the matter only as plagiarism remains conceptually inadequate because the more fundamental concern is the regulation of machine-assisted intellectual labor under conditions that can blur responsibility, distort evidential authorship, and compromise the rights of learners to fair and intelligible assessment standards.

4.8. Cross-Case Discussion Through the Theoretical Framework

Through a readiness lens, AUCA currently appears more publicly structured and ALU more publicly catalytic. Through a Technology Acceptance Model lens, both likely face strong adoption pressure because AI tools are widely perceived as useful and easy to use, but ALU's public communication more visibly embraces that attractiveness. Through Diffusion of Innovations, ALU's public discourse suggests higher observability and symbolic leadership, while AUCA's documentation suggests stronger compatibility work within internal systems and norms. The

comparative implication is that AI governance maturity is not one-dimensional. An institution may lead in innovation signaling but lag in public rule specificity, or vice versa. This distinction matters because it moves the analysis beyond a binary of adoption versus non-adoption and instead focuses on the alignment, or misalignment, between visible experimentation and visible governance. The more revealing distinction is between visible experimentation, visible governance, and visible alignment between them.

4.9. Implications for Rwandan Higher Education and Comparable African Contexts.

The key policy implication is that Rwanda's higher education sector now needs institution-level translation of national AI ambition into enforceable, pedagogically specific campus rules. The HEC has already created a sector benchmark (HEC, 2025). Universities now need locally legitimate versions of that benchmark, including institution-wide definitions of acceptable and unacceptable AI use, disclosure templates, staff-development pathways, assessment redesign principles, and privacy and data-governance safeguards. For comparable African contexts, the paper's broader implication is that universities do not need to wait for perfect infrastructure before governing AI well. What they do need is coherence. Governance credibility grows when institutions align mission, ethics, support systems, digital platforms, and classroom norms.

5. Conclusion and Recommendations

5.1. Summary of Core Findings

This study set out to compare publicly documented generative AI integration at ALU and AUCA. It found that both universities are visibly engaged with AI, but not in the same way. ALU's public presence is marked by AI research ambition, training initiatives, and academic-support programming. AUCA's public presence is marked by denser governance documents, digital-learning infrastructure, academic-integrity language, and AI-and-big-data institutionalization. Yet neither case currently presents a fully articulated, publicly accessible, institution-specific generative AI acceptable-use architecture at the level implied by Rwanda's newer national and sectoral policy documents. The claims advanced here concern visible governance and public institutional self-presentation rather than the full reality of internal practice.

5.2. Theoretical Contributions

The article contributes theoretically by distinguishing between innovation signaling and governance readiness and by showing how readiness, acceptance pressures, and diffusion dynamics interact rather than collapse into one another. Institutional AI maturity is not adequately captured by visible AI activity alone. It is better understood as an alignment problem among policy clarity, support systems, digital infrastructure, academic-integrity architecture, and educational purpose. The recommendations below follow directly from four findings established in the documentary record: visible AI activity at both universities, asymmetry in public governance density, weak public specificity on acceptable use, and incomplete visible alignment with the 2025 sectoral benchmark.

5.3. Policy and Institutional Recommendations

First, both universities should publish an institution-specific generative AI framework that moves beyond generic ethics or legacy plagiarism language. Second, the framework should differentiate assessment contexts, ideally using a permission model analogous to the HEC's red-amber-green logic. Third, both universities should require transparent AI-use declarations in student work and research outputs where machine assistance is relevant. Fourth, institution-wide faculty development should be formalized around assessment redesign, verification strategies, and discipline-specific AI pedagogy. Fifth, student AI literacy should include not only use skills but epistemic and ethical competencies: verification, bias recognition, source checking, privacy, and

truthful disclosure. These recommendations are not abstract ideals; they follow directly from the mismatch between visible AI activity and incomplete public governance in the cases studied.

5.4. Methodological Contributions

Methodologically, the article shows that document analysis can produce a serious, comparative first-stage study when the object of inquiry is formal institutional position, public communication, and visible governance architecture. It also demonstrates how secondary-data analysis can be used productively in fast-moving policy domains where primary data collection may lag behind institutional change.

5.5. Limitations

The study's main limitation is that it cannot show how staff and students actually interpret or enact these documents. Public guidance may overstate, understate, or simply fail to capture lived practice. Moreover, public websites are dynamic and may omit internal guidance unavailable to external researchers. The findings should therefore be read as an analysis of visible governance, not total governance.

5.6. Directions for Future Primary Research

A follow-up primary study should investigate at least six issues: stakeholder awareness of AI-related rules; actual patterns of generative AI use; differences across disciplines and assessment types; faculty confidence in managing AI-assisted work; student perceptions of fairness and disclosure; and the relationship between public policy architecture and lived practice. A mixed-method design combining surveys, interviews, and syllabus analysis would be especially suitable.

5.7. Concluding Reflection

The deeper significance of this study is philosophical as well as administrative. Universities are not simply deciding whether to “allow” or “ban” a tool. They are deciding what kind of intellectual formation they intend to defend in an age of machine-generated fluency. Rwanda's policy environment has already recognized that AI governance must be responsible, inclusive, and capacity-building. The challenge now is institutional translation. ALU and AUCA each possess real strengths. The central conclusion is not that one university has succeeded and the other failed. It is that both institutions currently occupy different positions within the same unresolved transition from AI visibility to institutionally articulated AI governance.

Conflicts of Interest: The authors declare no conflict of interest.

Reference

- Adventist University of Central Africa. (2021a). Online teaching and learning policy. <https://auca.ac.rw/wp-content/uploads/2024/02/Online-Policy-final-February-2021.pdf>
- Adventist University of Central Africa. (2021b). Student handbook 2018–2021. <https://auca.ac.rw/wp-content/uploads/2021/07/AUCA-STUDENT-HANDBOOK-2018-2021.pdf>
- Adventist University of Central Africa. (2022). AUCA research manual revised June 2022. <https://auca.ac.rw/wp-content/uploads/2022/06/AUCA-Research-Manual-Revised-June-2022.pdf>
- Adventist University of Central Africa. (2023). AUCA academic rules and regulations. <https://auca.ac.rw/wp-content/uploads/2024/02/Final-AUCA-Academic-Regulations-September-2023.pdf>
- Adventist University of Central Africa. (n.d.-a). Adventist University of Central Africa. <https://auca.ac.rw/>
- Adventist University of Central Africa. (n.d.-b). AUCA e-learning. <https://elearning.auca.ac.rw/eschool/>
- Adventist University of Central Africa. (n.d.-c). AUCA ICT policy. <https://auca.ac.rw/wp-content/uploads/2024/02/Final-AUCA-ICT-Policy-with-Cover.pdf>
- Adventist University of Central Africa. (n.d.-d). Research center. <https://web.auca.ac.rw/research>

- Adventist University of Central Africa. (n.d.-e). E-learning platform. <https://web.auca.ac.rw/academics/e-learning>
- African Leadership University. (n.d.-a). ALU policies. <https://www.alueducation.com/alupolicies/>
- African Leadership University. (n.d.-b). About the AI in Africa research project. <https://aiinafricaresearch.alueducation.com/about/>
- African Leadership University. (n.d.-c). Explore our global learning experiences. <https://www.alueducation.com/global-learning-experiences/>
- African Leadership University. (n.d.-d). Knowledge management. <https://library.alueducation.com/home>
- African Leadership University. (n.d.-e). Knowledge management: Upcoming events. <https://library.alueducation.com/upcoming-events>
- African Leadership University. (n.d.-f). BSc. (Hons) software engineering. <https://www.alueducation.com/bsc-hons-software-engineering/>
- An, Y., Yu, J. H., & James, S. (2025). Investigating the higher education institutions' guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. *International Journal of Educational Technology in Higher Education*, 22, Article 10. <https://doi.org/10.1186/s41239-025-00507-3>
- Belkina, M., Daniel, S., Nikolic, S., Haque, R., Lyden, S., Neal, P. R., Grundy, S., & Hassan, G. M. (2025). Implementing generative AI in higher education: A systematic review of case studies. *Computers and Education: Artificial Intelligence*, 8, 100407. <https://doi.org/10.1016/j.caeai.2025.100407>
- Bittle, K., & El-Gayar, O. (2025). Generative AI and academic integrity in higher education: A systematic review and research agenda. *Information*, 16(4), 296. <https://doi.org/10.3390/info16040296>
- Bond, M., Khosravi, H., De Laat, M., Bergdahl, N., Negrea, V., Oxley, E., Pham, P., Chong, S. W., & Siemens, G. (2024). A meta systematic review of artificial intelligence in higher education: A call for increased ethics, collaboration, and rigour. *International Journal of Educational Technology in Higher Education*, 21, Article 4. <https://doi.org/10.1186/s41239-023-00436-z>
- Bowen, G. A. (2009). Document analysis as a qualitative research method. *Qualitative Research Journal*, 9(2), 27–40. <https://doi.org/10.3316/QRJ0902027>
- Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, 20, Article 38. <https://doi.org/10.1186/s41239-023-00408-3>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M., Al-Busaidi, K. A., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., ... Wright, R. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
- Higher Education Council. (2025). Guidelines and assessment tools for the virtual learning and the use of artificial intelligence in Rwanda's higher learning institutions. https://www.hec.gov.rw/fileadmin/user_upload/HEC/Publications/Guidelines/GUIDELINES_ASSESMENT_TOOLS_FOR_VL_and_AI_in_HLIs.pdf
- Hsieh, H.-F., & Shannon, S. E. (2005). Three approaches to qualitative content analysis. *Qualitative Health Research*, 15(9), 1277–1288. <https://doi.org/10.1177/1049732305276687>
- Ifenthaler, D., & Yau, J. Y.-K. (2020). Utilising learning analytics to support study success in higher education: A systematic review. *Educational Technology Research and Development*, 68, 1961–1990. <https://doi.org/10.1007/s11423-020-09788-z>
- Jensen, L. X., Buhl, A., Sharma, A., & Bearman, M. (2025). Generative AI and higher education: A review of claims from the first months of ChatGPT. *Higher Education*, 89, 1145–1161. <https://doi.org/10.1007/s10734-024-01265-3>

- Jin, Y., Yan, L., Echeverria, V., Gašević, D., & Martinez-Maldonado, R. (2025). Generative AI in higher education: A global perspective of institutional adoption policies and guidelines. *Computers and Education: Artificial Intelligence*, 8, 100348. <https://doi.org/10.1016/j.caeai.2024.100348>
- Johnston, M. P. (2014). Secondary data analysis: A method of which the time has come. *Qualitative and Quantitative Methods in Libraries*, 3(3), 619–626. <https://www.qqml-journal.net/index.php/qqml/article/view/169>
- Kasneji, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneji, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kim, J., Klopfer, M., Grohs, J. R., Eldardiry, H., Weichert, J., Cox, L. A., II, & Pike, D. (2025). Examining faculty and student perceptions of generative AI in university courses. *Innovative Higher Education*, 50(4), 1281–1313. <https://doi.org/10.1007/s10755-024-09774-w>
- Lodge, J. M., Thompson, K., & Corrin, L. (2023). Mapping out a research agenda for generative artificial intelligence in tertiary education. *Australasian Journal of Educational Technology*, 39(1), 1–8. <https://doi.org/10.14742/ajet.8695>
- Luo, J. (2024). A critical review of GenAI policies in higher education assessment: A call to reconsider the “originality” of students’ work. *Assessment & Evaluation in Higher Education*, 49(5), 651–664. <https://doi.org/10.1080/02602938.2024.2309963>
- Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000386693>
- Ministry of ICT and Innovation [MINICT]. (2023). The National AI Policy for the Republic of Rwanda. Government of Rwanda. https://www.minict.gov.rw/fileadmin/user_upload/minict_user_upload/Documents/Policies/Artificial_Intelligence_Policy.pdf
- Moorhouse, B. L., Yeo, M. A., & Wan, Y. (2023). Generative AI tools and assessment: Guidelines of the world’s top-ranking universities. *Computers and Education Open*, 5, 100151. <https://doi.org/10.1016/j.cao.2023.100151>
- OECD. (2023). OECD digital education outlook 2023: Towards an effective digital education ecosystem. OECD Publishing. <https://doi.org/10.1787/c74f03de-en>
- One Million Rwandan Coders. (n.d.). Our digital future. <https://uat.onemillioncoders.rw/about>
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press. <https://teddykw2.wordpress.com/wp-content/uploads/2012/07/everett-m-rogers-diffusion-of-innovations.pdf>
- Sangwa, S. ., & Mutabazi, P. (2025). Generative AI and Academic Integrity in Online and Distance Learning: A Policy Index and Evidence for Assessment Redesign in the Global South. *Open Journal of AI, Ethics & Society* (ISSN: 3105-3076), 1(2). <https://doi.org/10.65655/dg1rsd17>
- Sebihi, A., Schoelen, L., & Uwamwezi, B. (2025). Unveiling AI’s impact in Rwandan higher education: Navigating trends, challenges, and imperatives at the University of Rwanda. *Discover Education*, 4, Article 166. <https://doi.org/10.1007/s44217-025-00576-6>
- Theodorio, A. O., Waghid, Z., Mataka, T. W., & Adegoke, O. (2024). Demystifying Lesotho, Rwandan and Nigerian educators’ viewpoints on smart technologies supporting AI in higher education. *Education and Information Technologies*, 29, 24285–24307. <https://doi.org/10.1007/s10639-024-12820-x>
- UNESCO. (2021). Recommendation on the ethics of artificial intelligence. <https://www.unesco.org/en/articles/recommendation-ethics-artificial-intelligence>
- UNESCO. (2022). Recommendation on the ethics of artificial intelligence. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>
- UNESCO. (2024). AI competency framework for teachers. UNESCO. <https://www.unesco.org/en/articles/ai-competency-framework-teachers>
- UNESCO. (2025). AI and education: Protecting the rights of learners. UNESCO. <https://www.unesco.org/en/articles/ai-and-education-protecting-rights-learners?hub=195885>

University of KwaZulu-Natal. (2025, January). Principles and guidelines on the use of artificial intelligence tools for academic work. <https://ukzn.ac.za/wp-content/uploads/2025/01/UKZN-AI-Academic-Guidelines-SE-01-0312-24-protected.pdf>

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.