

AI Literacy and Institutional Governance Response among Leaders in Private Higher Education

Innocent Kumwenda*, Joe Likando., Misheck Malama., Happy Siame., Davy Mainde

Africa Research University, Lusaka, Zambia

*Correspondence Author

DOI: <https://doi.org/10.47772/IJRISS.2026.100800837>

Received: 30 August 2026; Accepted: 04 September 2026; Published: 18 September 2026

ABSTRACT

Artificial intelligence tools are already in daily use across Zambia's private higher education institutions, yet little is known about whether the leaders responsible for governing that use are equipped to do so, or whether their institutions have built the structures to support them. This paper reports findings from a sequential exploratory mixed-methods study examining governance-level AI literacy and institutional AI response among strategic leaders across Zambia's twelve Tier 1 private universities. A structured questionnaire completed by one hundred governance-role staff was combined with semi-structured interviews conducted with twelve strategic leaders purposively drawn from five institutions, and the two strands were integrated using the Technology Acceptance Model as an interpretive lens. The study finds that governance-level AI literacy, the capacity to evaluate, question, and be accountable for AI's institutional use rather than simply operate the tools, is markedly underdeveloped among leaders regardless of seniority or academic qualification, and that this deficit is mirrored, and compounded, by an almost total absence of formal institutional AI governance: defined strategies, adoption procedures, and ethical safeguards are the exception rather than the rule. Leaders consistently recognise AI's strategic importance in principle while lacking the specific capacity to govern it in practice, a gap this paper argues is structural rather than personal, rooted in professional formation systems that did not anticipate AI governance as a leadership competency. The findings carry direct implications for institutional capacity-building, sector regulation, and national policy, and this paper's originality lies in providing the first empirically grounded account of the AI-literacy-to-governance pathway, and its breakdown, within Africa's private higher education sector.

Keywords: AI literacy; institutional governance; strategic leadership; higher education; Zambia

INTRODUCTION

Artificial intelligence has entered African higher education faster than the governance structures meant to oversee it. Students and staff at Zambia's private universities now use generative AI tools routinely, for drafting, research, assessment preparation, and administrative work, yet this diffusion has occurred largely outside any formal institutional decision about whether, how, or under what safeguards such use should proceed. The strategic leaders nominally responsible for that oversight, Vice Chancellors, Deputy Vice Chancellors, Registrars, Deans, and Heads of Department, occupy an unusual position: they are expected to govern a technology whose technical workings, ethical risks, and vendor claims most of them have had no formal opportunity to learn to evaluate.

This gap between adoption and governance capacity is not unique to Zambia, but it carries particular weight in a private higher education sector defined by fee dependency, resource constraint, and regulatory ambiguity. Where public universities may draw on ministry guidance, national research infrastructure, or sector-wide standards, Zambia's private institutions largely govern themselves, competing for the same fee-paying student base while carrying full responsibility for the pedagogical and ethical consequences of the technology their students and staff already use. Whether these institutions' leaders possess the literacy required to discharge that

responsibility, and whether their institutions have built the governance structures to support them in doing so, has not previously been established empirically.

The literature offers two well-developed but separate bodies of theory relevant to this question. The Technology Acceptance Model and its extensions explain how perceived usefulness and perceived ease of use shape individual adoption behaviour, providing a lens for understanding why leaders might embrace or resist engaging with AI governance as a task. A parallel and more recent literature on AI literacy, most influentially UNESCO's framework distinguishing basic AI awareness from the higher-order capacity for accountable AI governance, specifies what leaders would need to know and be able to do to govern responsibly, as opposed to merely use AI tools themselves. Neither body of theory, however, has been tested empirically against the specific conditions of African private higher education, where governance decisions are made under simultaneous financial, infrastructural, and regulatory constraint.

The distinction between using AI and governing AI is not a semantic nicety; it is the conceptual hinge on which this paper's entire argument turns. A leader can be a confident, even enthusiastic, personal user of generative AI tools while remaining wholly unequipped to discharge the governance obligations institutional deployment creates: assessing a vendor's data-handling practices, judging whether an assessment tool produces equitable outcomes across a diverse student body, or deciding what safeguards a policy ought to contain before a tool is approved for use. AI literacy frameworks that collapse this distinction risk certifying leaders as literate on the basis of comfort with the technology rather than capacity to govern it, which is precisely the higher-order competency this paper investigates. Regionally, the evidence base for how this distinction plays out in practice remains thin. Studies of AI adoption in African higher education have tended to document usage patterns, student and staff attitudes, or infrastructural readiness, leaving the governance capacity of institutional leadership comparatively under-examined. Recent scholarship has begun to catalogue AI governance frameworks and algorithmic-accountability practice across African higher education more broadly (Memarian and Doleck, 2023; Sangwa et al., 2026; Jacobs-Basadien and Pather, 2026), and closely related governance-readiness evidence has begun to emerge from other Sub-Saharan university settings (Dusengumuremyi et al., 2026), but this literature has not yet been brought to bear on the specific conditions of Zambia's private HEI sector, a gap this paper's focus is intended to help close.

This paper addresses that gap directly. Drawing on the first of four objectives from a wider doctoral investigation into how strategic leaders in Zambia's private higher education sector are adapting to AI-driven change, this paper asks two linked empirical questions: what is the current level of governance-level AI literacy among strategic leaders across Zambia's Tier 1 private universities, and how, if at all, has that literacy translated into formal institutional AI governance. Investigating literacy and institutional response together, rather than in isolation, matters because the relationship between the two is precisely where the practical governance problem lies. A leader who privately recognises AI's importance but lacks the specific competence to evaluate a vendor's claims, or an institution that has not yet built a committee, policy, or procedure through which that evaluation could occur, each represents a different point of failure in the same governance chain, and a coherent capacity-building response requires knowing where along that chain the failure is concentrated.

The paper proceeds as follows. Section 2 sets out the sequential exploratory mixed-methods design used to investigate both literacy and institutional response. Section 3 presents the findings, first for individual AI literacy, then for institutional AI response, before examining how the two interact. Section 4 concludes by arguing that the observed gap between individual recognition and institutional action is a governance capacity problem rather than a motivational one, and sets out the implications of that argument for institutional leaders, the sector regulator, and national policymakers.

METHODOLOGY

Research Design

This paper draws on data collected for a wider sequential exploratory mixed-methods investigation into how strategic leaders across Zambia's private higher education sector are adapting to AI-driven change. The design followed a two-phase QUAL → quan sequence: an initial qualitative phase generated the contextual

understanding needed to build a locally grounded quantitative instrument, and the two strands were subsequently integrated at the interpretation stage. This design was selected because AI governance practice among Zambian private HEI leaders had not previously been documented in sufficient depth to support a purely quantitative instrument, and because the study's broader aim of developing a contextual governance model required qualitative depth that a quantitative-only design could not supply.

Population and Sampling

The target population comprised senior academic and administrative staff holding governance responsibility across the twelve Tier 1 private universities registered by the Higher Education Authority of Zambia, an estimated population of 156 individuals spanning five role categories: Vice Chancellor/Principal, Deputy Vice Chancellor, Registrar, Dean of School, and Head of Department. For the quantitative phase, a sample of 113 was calculated using Slovin's formula at a 95 percent confidence level and drawn through stratified random sampling by role category across all twelve institutions; 100 usable responses were returned, an 88.5 percent response rate. For the qualitative phase, twelve strategic leaders were purposively selected, using maximum variation sampling to capture diversity in institutional ownership type, size, and operating history, from five of the twelve institutions, spanning all five role categories. Interviews ran between 46 and 68 minutes, were audio-recorded, and were transcribed word-for-word before being shared back with each participant to confirm accuracy.

Restricting the sampling frame to Tier 1 institutions, rather than the private HEI sector as a whole, was a deliberate analytical choice. Tier 1 status denotes the highest classification level within Zambia's higher education qualifications framework, and the leaders who govern these institutions carry the most consequential and most complex AI-related decision-making authority within the private sector; if governance-level literacy and institutional structure are found wanting even at this most institutionally mature tier, the implication for less-resourced tiers of the sector is unlikely to be more favourable. The stratified probability sample underlying the quantitative phase supports descriptive generalisation of the literacy and institutional-response findings across the twelve-institution Tier 1 population, subject to ordinary sampling error; the qualitative findings, drawn from five of the twelve institutions, are reported as illustrative and explanatory rather than as representative of the full sector, a distinction maintained throughout the presentation of results below.

Measures and Instrumentation

The findings reported here draw on two sections of the study's data collection instruments. Governance-level AI literacy was measured using six Likert-scale items covering conceptual understanding of AI, critical evaluation capacity, ethics and governance awareness, confidence leading AI-related strategic discussions, the ability to distinguish vendor claims from verified performance, and awareness that AI systems can produce unequal outcomes across demographic groups. Institutional AI response was measured using a parallel set of six items addressing the presence of a defined AI strategy, active leadership promotion of AI-driven innovation, alignment between AI integration and institutional mission, ethical grounding of AI adoption, the existence of formal tool-evaluation procedures, and the degree to which AI use across the institution was driven by individual initiative rather than coordinated policy. Both item sets were developed from themes generated in the qualitative phase and returned excellent internal consistency, with Cronbach's alpha coefficients of 0.907 for the literacy items and 0.932 for the institutional response items, comfortably exceeding the conventional 0.70 threshold for social science instruments. The corresponding sections of the semi-structured interview guide explored the same two constructs in depth, allowing the qualitative accounts to explain and contextualise the quantitative patterns.

Data Analysis

Quantitative data were analysed descriptively, through frequencies, percentages, means, and standard deviations, consistent with the study's exploratory rather than confirmatory analytical purpose at this stage. Composite-level inferential testing was additionally conducted to check whether the AI Literacy pattern reported in Section 3.2 varied systematically by respondent role, tenure, qualification, or institutional size, and whether the AI Literacy and Institutional AI Response composites were statistically associated with one another; because Shapiro-Wilk tests indicated the composite scores departed from a normal distribution, Kruskal-Wallis H tests were used in place of one-way ANOVA for the group comparisons, and Pearson correlation, robustness-checked against

Spearman's rank correlation, for the bivariate association, each assessed at a significance threshold of $\alpha = 0.05$. These results are reported in Section 3.5. Qualitative data were analysed using Braun and Clarke's reflexive six-phase thematic analysis procedure, with a member of the supervision team independently coding a subsample of transcripts to support analytical dependability. Quantitative and qualitative findings were integrated at the interpretation stage, with the qualitative accounts used to explain the mechanisms underlying the patterns the quantitative data revealed, consistent with the analytical priority the sequential exploratory design assigns to the qualitative strand. Interview participants are identified throughout by anonymised codes running from P1 through P12, tagged only with a role label; no institution is named alongside any individual account.

Methodological Justification

Several design choices warrant brief methodological justification. A mixed-methods approach was preferred over either a purely quantitative or purely qualitative alternative because the research problem is simultaneously measurable, the prevalence of literacy and governance gaps across a twelve-institution sector, and interpretively complex, the mechanisms by which those gaps arise and persist within specific institutional cultures; a single-method design would have captured one dimension at the expense of the other. The qualitative-first sequencing reflected the absence of any prior empirically grounded account of AI governance practice in this sector: developing a quantitative instrument without first understanding, qualitatively, what governance-relevant constructs actually mattered to leaders in this specific context risked measuring variables imported from unrelated settings rather than the sector's own lived governance realities. Where the two strands agreed, interview accounts independently echoing the same patterns the questionnaire had already measured, that convergence is read as corroborating evidence, lending the overall findings more weight than either strand could carry on its own, though it is not a substitute for the causal or moderation testing a larger, purpose-built quantitative study would need to support.

Validity Considerations

One further validity consideration concerns the plausibility of social desirability bias, given that respondents were reporting on their own, and their institution's, governance shortcomings. Three features of the data make this an unlikely explanation for the pattern observed. First, any such bias would be expected to operate in the direction of understating rather than exaggerating governance gaps, meaning the true extent of the deficit is, if anything, more likely to be understated than overstated by these findings. Second, participants offered detailed, unflattering admissions on their own initiative, most notably that AI tools already in daily use at their institutions had never been through any kind of vendor evaluation, a disclosure a self-protective account would be unlikely to include. Third, when interview transcripts were sent back to participants to check, none of these unflattering passages came back amended or withdrawn, which suggests participants stood by the picture they had given rather than having overstated it in the moment.

RESULTS

Who Was Surveyed and Interviewed

The 100 questionnaire respondents were distributed across the five governance role categories as follows: Heads of Department, 58 percent; Deans of School, 28 percent; Deputy Vice Chancellors, 5 percent; Registrars, 5 percent; and Vice Chancellors, 4 percent, a distribution broadly consistent with the relative size of each role stratum within Tier 1 institutional leadership structures. Respondents were experienced by tenure: 84 percent had held a strategic leadership role for five years or more, split evenly between the five-to-ten-year and eleven-to-fifteen-year bands (31 percent each), with a further 22 percent reporting more than fifteen years in post and 16 percent reporting less than five years. Formal academic attainment was high, with 85 percent holding a postgraduate qualification at Master's (44 percent) or Doctoral (41 percent) level, and the remaining 15 percent holding another professional or higher diploma qualification. Institutional size among respondents skewed toward the sector's medium-sized institutions (51 percent), with smaller proportions from the smallest (26 percent) and largest (23 percent) Tier 1 universities. The twelve interview participants spanned all five role categories, including two Vice Chancellors, two Deputy Vice Chancellors, two Registrars, two Deans of School, and four Heads of Department, drawn from five institutions chosen to capture variety across the sector in who

owns them and how long each has been operating. This profile matters for what follows: the leaders reporting low AI literacy in Section 3.2 are not a junior or inexperienced cohort. They are, on the whole, senior, tenured, and highly credentialed, which sharpens rather than softens the significance of the literacy gap the data reveal. A competing explanation for weak AI literacy scores, that respondents were simply new to their roles or early in their careers and had not yet had the opportunity to develop relevant expertise, finds little support in this profile: leaders with over a decade in strategic leadership, and leaders holding doctoral qualifications, reported governance-level AI literacy no higher, on the whole, than their less experienced or less credentialed colleagues, a pattern more consistent with a gap in what professional formation and career development have covered than with a gap that ordinary career progression would be expected to close on its own.

A Governance-Level AI Literacy Deficit That Cuts Across Seniority

Every one of the six AI literacy items returned a mean score below the neutral midpoint of the five-point scale, indicating that disagreement, not agreement, was the dominant response across the sample (see Table 1). Leaders rated themselves comparatively higher on general conceptual understanding of AI than on any of the other five items, but even that highest-scoring item fell short of neutral agreement. Scores declined further, in a fairly consistent gradient, as items moved from general awareness toward the specific accountable judgements that institutional governance actually requires: critically evaluating an AI tool for institutional use, understanding the ethical and governance issues AI raises, feeling confident enough to lead a strategic discussion about it, recognising that AI systems can behave differently across demographic groups, and, scoring lowest of all six items, distinguishing a vendor's claims about an AI product from independently verified evidence of how it actually performs.

That gradient is not a statistical curiosity; it maps directly onto the distinction between describing what AI is, which most respondents can do with some confidence, and being accountable for a specific governance decision about it, which most cannot. Interview accounts gave this pattern texture and specificity. One Vice Chancellor participant, asked to appraise his own literacy against what institutional governance actually demands, distinguished sharply between the two: he could describe in general terms what generative AI tools do, but if asked to judge whether a particular tool proposed for student assessment was fair, transparent, and produced equitable outcomes across the student body, he said plainly that he would not know where to begin, even though that judgement was precisely the kind of decision his office might be called on to make. A second Vice Chancellor participant located the difficulty less in personal limitation than in a structural feature of under-resourced private institutions: without a technically literate advisor able to translate AI's complexity into governance-relevant terms, he found himself expected to be simultaneously the person who understood the technology and the person who decided what to do about it, two demanding roles neither his training nor his institution's resourcing had prepared him to hold at once.

The lowest-scoring item, the ability to distinguish a vendor's claims from independently verified performance, deserves particular attention because it sits at the exact point where AI literacy becomes consequential rather than merely reflective. A leader who cannot evaluate a vendor's claims is not simply under-informed; they are structurally exposed to whatever a commercial AI provider chooses to represent about a product's accuracy, fairness, or data-handling practices, with no independent basis on which to interrogate that representation before committing institutional resources or student data to it. This is precisely the capacity Miao et al.'s UNESCO framework identifies as the highest-order tier of AI literacy, the tier concerned with governance and accountability rather than conceptual awareness or basic operational competence, and it is also the tier at which this study's respondents scored weakest of all.

A further account, from a Registrar participant, illustrated how an institution's very entry point into AI governance can shape everything that follows. AI only reached his institution's leadership agenda after examiners started catching AI-generated submissions and treated it as an academic integrity emergency; there had been no earlier, calmer moment of deliberate strategic planning through which the issue might otherwise have surfaced. Once governance attention is captured by AI's most alarming manifestation rather than its potential value, reorienting that framing toward a more constructive, enabling posture proves difficult, a dynamic visible across several of the institutional response findings reported in Section 3.3. High formal qualification did not appear to close the literacy gap: doctoral study, as one Head of Department participant observed of his own postgraduate

training, had not included a single module addressing artificial intelligence, algorithmic systems, or data governance, leaving him well prepared to lead an academic department but not to judge whether a vendor's AI assessment tool handled student data responsibly. Eight of the twelve interview participants reported never having received formal training in AI ethics or governance of any kind. Taken together, the quantitative and qualitative evidence point to a deficit with roots in the formation system rather than in any individual leader: doctoral programmes and executive development pathways alike were built before AI governance existed as a recognisable competency to teach, and no amount of individual effort or accumulated professional experience substitutes for training that was simply never offered.

Table 1

Table 1. Mean scores and reliability for AI Literacy (Section B) and Institutional AI Response (Section C) items (n = 100)

Construct	Item	Mean	SD
AI Literacy ($\alpha = 0.907$)	Understands AI concepts relevant to HE leadership	2.92	1.12
AI Literacy	Can critically evaluate AI tools for institutional application	2.66	1.19
AI Literacy	Understands ethical and governance issues related to AI	2.62	1.13
AI Literacy	Feels confident leading AI-related strategic discussions	2.55	1.07
AI Literacy	Aware AI tools may produce unequal outcomes across demographic groups	2.51	1.15
AI Literacy	Can distinguish vendor claims from independently verified performance	2.44	1.09
Institutional AI Response ($\alpha = 0.932$)	AI use is driven by individual initiative rather than coordinated policy	3.77	1.14
Institutional AI Response	AI integration aligns with institutional mission	2.92	1.11
Institutional AI Response	Leadership actively promotes AI-driven innovation	2.70	1.15
Institutional AI Response	AI adoption is guided by ethical standards	2.54	1.10
Institutional AI Response	Institution has formal procedures to evaluate AI tools before adoption	2.31	1.13
Institutional AI Response	Institution has a defined AI-related strategy	2.27	1.08

Note. Five-point Likert scale (1 = Strongly Disagree to 5 = Strongly Agree). Items ordered by descending mean within each construct, except the reverse-signalling item, which is listed first for its construct as its high mean indicates governance absence rather than governance strength.

Institutional AI Governance Is the Exception, Not the Rule

If individual literacy was weak, formal institutional response was weaker still. Of the six institutional response items, five fell below the neutral midpoint, and the one exception, on closer inspection, actually confirms the same absence rather than contradicting it: it captured the extent to which AI adoption was happening through unplanned staff-level initiative rather than any coordinated institutional plan, so a high score here signals ungoverned, ad hoc use rather than a deliberate institutional stance. The lowest-scoring item in the entire questionnaire, across all sections, concerned whether the respondent's institution had a defined AI-related

strategy; a clear majority disagreed or strongly disagreed that one existed, and only a small minority agreed that it did. A closely related item, on whether formal procedures existed to evaluate an AI tool before it was adopted, returned an almost identical pattern. Only a small minority of respondents agreed that AI adoption at their institution was guided by ethical standards at all.

The interview accounts described what this absence looks like operationally, not merely as a survey response but as daily institutional practice. The Registrar participant quoted in Section 3.2 gave a precise inventory of the vacuum at his own institution: there was no dedicated AI policy, nobody with formal oversight of AI-related decisions, and nothing resembling a check before a new tool entered an academic programme; the closest document on file was a general acceptable-use policy for technology, written years before generative AI tools reached the public. AI, he noted, had been in continuous daily use at the institution since that point, while the institution's governance documents continued to treat it as though it did not exist. A Deputy Vice Chancellor participant described a related but distinct governance failure, one of institutional bandwidth rather than institutional will: he had repeatedly attempted to place an AI governance policy on the institution's Senate agenda, only to see it deferred each time in favour of more urgent business, business that, he observed with some irony, frequently consisted of academic integrity cases arising from exactly the kind of ungoverned AI use a policy would have addressed. The institution was, in effect, absorbing the costs of absent AI governance while lacking the time that developing such governance would require, a self-perpetuating pattern rather than a simple oversight.

The weak showing on ethical grounding of AI adoption carries particular weight given the populations these institutions serve. Many of Zambia's private HEIs enrol substantial numbers of first-generation university students from lower-income backgrounds, precisely the demographic for whom algorithmic bias and inequitable AI-driven outcomes, documented in the wider literature on AI fairness and algorithmic accountability (Buolamwini and Gebru, 2018; Pasipamire and Muroyiwa, 2024), carry the greatest and least visible cost. Absent ethical AI governance is not, for this population, an abstract policy gap; it is a live equity exposure operating without the institutional safeguards that would ordinarily be expected to catch it.

It is worth being precise about what the institutional response data can and cannot support. The items reported here describe respondents' perceptions of their own institution's governance posture; they were not designed, and are not used here, to test a causal relationship between the absence of formal AI governance and any specific downstream harm. What the data do establish, with reasonable confidence given the size and coverage of the sample, is the prevalence of the governance vacuum itself: across a stratified sample spanning all twelve Tier 1 institutions and all five governance role categories, the pattern of absent strategy, absent evaluation procedure, and weak ethical grounding was remarkably consistent rather than concentrated in a handful of outlying institutions. That consistency is itself a finding of some weight, since it suggests a sector-wide condition rather than a scattering of isolated institutional failures.

A Self-Reinforcing Gap Between Recognition and Action

Read together, the literacy and institutional response findings describe a single governance chain breaking at two connected points rather than two unrelated problems. Every interview participant, without exception, affirmed AI's strategic importance to their institution's future; none treated it as a peripheral concern. Yet that near-universal recognition did not translate into institutional governance action, and the data suggest why: leaders who have not themselves been equipped with governance-level AI literacy are poorly positioned to build the institutional structures, policies, and evaluation procedures that governance requires, while institutions without such structures in turn provide no mechanism through which leaders' literacy could be developed, tested, or supported in practice. Each condition sustains the other. Read through the lens of the Technology Acceptance Model, this pattern is consistent with a workforce that perceives AI's usefulness clearly but experiences AI governance itself, as opposed to AI use, as neither easy nor within reach; the questionnaire did not administer validated Technology Acceptance Model scales, so this connection is offered as an interpretive frame for the pattern observed rather than as a direct empirical test of the model's constructs.

This self-reinforcing pattern is also consistent with organisational readiness theory, which holds that meaningful organisational change requires both change commitment, the collective belief that change matters, and change

efficacy, the collective belief that the organisation is capable of implementing it (Weiner, 2009). Every interview participant supplied the first ingredient unprompted; almost none could point to evidence of the second. Institutional theory offers a complementary explanation for why this gap has persisted so uniformly across a twelve-institution sector rather than resolving unevenly, as idiosyncratic institutional differences alone might predict: in the absence of coercive pressure from a sector regulator, normative pressure from a professional body, or a visible model to imitate among peer institutions, organisations facing a novel governance demand tend to converge on inaction rather than independently developing divergent responses (DiMaggio and Powell, 1983; Meyer and Rowan, 1977). At the time of writing, the Higher Education Authority of Zambia had issued no AI-specific compliance requirement, no professional association had issued AI governance guidance for the sector, and none of the institutions studied here had yet built a governance framework visible enough for peer institutions to observe and imitate. Each of the three isomorphic pressures institutional theory identifies was, in other words, absent, offering one explanation, beyond leaders' own literacy limitations, for why formal AI governance failed to emerge even where individual leaders recognised its importance.

The practical significance of this framing is that it points toward a different intervention than either finding would suggest in isolation. Addressing the literacy deficit through individual training alone, without simultaneously building the institutional structures, in defined strategies, adoption procedures, and oversight committees, through which that literacy could be exercised, would leave newly literate leaders with nowhere to apply what they had learned. Conversely, mandating institutional AI governance frameworks without building the underlying leadership literacy to operate them risks producing policy documents that exist on paper but cannot be implemented with genuine understanding. The evidence instead points to a coordinated response: literacy development and institutional structure-building addressed together, and addressed as matters of urgency rather than as items competing, as the Deputy Vice Chancellor's account illustrated, for space on an already crowded governance agenda.

Checking the Literacy Pattern Against Demographic and Institutional Variation

Because Section 3.2's literacy gradient held so consistently across the sample, composite-level inferential tests were run to check whether the AI Literacy pattern could instead be explained by who happened to respond, against four candidate variables: role category, years in strategic leadership, highest academic qualification, and institutional size. Shapiro-Wilk tests indicated the AI Literacy composite departed from a normal distribution ($W = 0.970$, $p = 0.024$), so Kruskal-Wallis H tests were used in place of one-way ANOVA. AI Literacy did not differ significantly by role category ($H(4) = 5.88$, $p = .209$), years in strategic leadership ($H(3) = 2.01$, $p = .571$), or highest academic qualification ($H(2) = 0.57$, $p = .751$): the deficit documented in Section 3.2 was not confined to, or explained by, any particular seniority band, tenure group, or qualification level, reinforcing this paper's central claim that the gap is structural rather than personal. AI Literacy did, however, differ significantly by institutional size ($H(2) = 6.28$, $p = .043$), with respondents from larger institutions reporting somewhat higher literacy scores, on average, than those from smaller and medium-sized institutions, plausibly reflecting the wider vendor exposure, larger technical support base, and greater administrative capacity that larger institutions are more likely to command. This single significant result should be read cautiously given the moderate, unevenly sized cells the three size bands produce (26, 51, and 23 respondents respectively; Section 3.1) and is offered as a descriptive association rather than a causal claim.

A Pearson correlation between the AI Literacy and Institutional AI Response composites, robustness-checked against Spearman's rank correlation with no material change in the pattern, found no statistically significant association at the composite level ($r = -0.104$, $p = .303$). Read alongside the qualitative evidence in Sections 3.3 and 3.4, this null composite-level result should not be mistaken for evidence against the self-reinforcing relationship this paper describes: a single cross-sectional correlation coefficient is a blunt instrument for detecting a relationship that, as the interview accounts suggest, operates through specific institutional mechanisms, such as a Vice Chancellor's admitted uncertainty translating into a Senate agenda item repeatedly deferred, rather than as a general linear association running uniformly across all one hundred respondents. This result is reported here in the interest of transparency about what the quantitative strand can and cannot establish, and it sharpens rather than weakens this paper's reliance on the qualitative evidence to explain the mechanism connecting individual literacy to institutional inaction.

Two further comparisons this pattern might reasonably invite, direct testing across institutional ownership type (faith-based versus secular, for example) and a fully powered inferential comparison across all five leadership-role strata individually, were not possible with the present data. Institutional ownership type was used to guide the purposive selection of the five institutions included in the qualitative phase (Section 2) but was not captured as a variable in the quantitative instrument, and the smallest role strata (Vice Chancellors, $n = 4$; Deputy Vice Chancellors, $n = 5$; Registrars, $n = 5$) are too small to support a reliably powered statistical comparison. Capturing institutional ownership type directly in the quantitative instrument, and deliberately oversampling the senior-role strata, are accordingly identified here as specific design refinements for the larger-sample follow-up study called for in Section 4.

CONCLUSION

This paper set out to establish, empirically, how governance-level AI literacy and institutional AI governance stand among strategic leaders across Zambia's Tier 1 private higher education sector, and how the two relate to one another. The evidence is unambiguous on both counts. Governance-level AI literacy, the capacity to evaluate, question, and take accountability for AI's institutional use rather than simply describe or operate AI tools, is markedly underdeveloped, and that deficit holds regardless of seniority, tenure, or formal academic qualification. Formal institutional AI governance is, with rare exception, absent: defined strategies, adoption procedures, and ethical safeguards were the exception rather than the rule across the twelve-institution sector studied. Neither finding is explained by indifference. Every leader interviewed recognised AI's strategic importance; what was missing was the specific capacity, and the institutional infrastructure, to act on that recognition responsibly.

This paper's principal contribution is establishing that the literacy deficit and the governance vacuum are not separate problems but two expressions of a single, self-reinforcing gap, evidence that has not previously been documented for Africa's private higher education sector and that has direct implications for how institutions, regulators, and policymakers should respond. For institutional leaders, the priority is twofold and should not wait for external regulatory mandate: developing institutional AI governance frameworks, drawing on and adapting existing international guidance to local resource and regulatory conditions, alongside investing in AI literacy development that is contextually grounded rather than imported wholesale from resource environments unlike their own. Where an individual institution lacks the capacity to develop such literacy content independently, collaboration through existing sector associations offers a way of sharing both the cost and the expertise required, rather than each of the twelve institutions attempting to solve the same problem in isolation. The evidence bears on the Higher Education Authority of Zambia's role too: it supports the case for sector-specific AI governance standards, introduced first through a period of guidance and only later through compliance enforcement, so that institutions making genuine good-faith progress are not penalised for the transition they are still working through; folding AI governance readiness into existing quality assurance criteria would supply exactly the compliance-oriented push this sector's institutional cultures seem to need to move from aspiration to action. For national policymakers, the findings argue for a purpose-built policy framework for AI in higher education, one that speaks directly to the funding and infrastructure constraints distinctive to privately funded institutions rather than one calibrated to a better-resourced public or international setting, together with clearer sector-specific guidance on how existing data protection law applies to AI platforms operating in higher education.

To make the first of these recommendations actionable rather than aspirational, three concrete starting points are offered. First, a minimal governance structure need not be elaborate to be functional: a small AI governance committee, chaired by a Deputy Vice Chancellor or equivalent, with terms of reference covering tool approval, data-protection compliance, and academic-integrity implications, and reporting formally to Senate or its equivalent, converts the ad hoc, individually driven adoption documented in Section 3.3 into a traceable institutional decision. Second, a short vendor-evaluation checklist, covering data-handling and storage location, evidence of bias or fairness testing, explainability of outputs, cost and support terms, and a defined institutional exit route if the tool is discontinued, gives leaders a concrete instrument for the specific competency, distinguishing vendor claims from verified performance, that Section 3.2 found weakest. Third, literacy development need not begin as a bespoke curriculum: a short, contextually grounded module sequence covering

foundational AI concepts, the ethics and governance issues specific to higher education, a practicum applying the vendor-evaluation checklist to a real tool under consideration, and case discussion drawn from the kind of institutional dilemmas this study's participants described, offers institutions an accessible entry point that does not require importing training content built for a different resourcing context. None of these three starting points requires capacity beyond what a resource-constrained Tier 1 institution already possesses; what each requires is the decision to formalise practice that, at present, is simply not happening.

Although the empirical evidence reported here is specific to Zambia's Tier 1 private higher education sector, the underlying governance dynamic, individual recognition of a technology's strategic importance outpacing the institutional literacy and structure needed to govern it responsibly, is unlikely to be unique to this sector or this country. The institutional conditions driving the gap documented here, fee-dependent funding that places governance investment in competition with core operations, the absence of sector-specific regulatory guidance, and professional formation systems that did not anticipate AI governance as a leadership competency, are shared to varying degrees by private higher education systems elsewhere in Sub-Saharan Africa, and by other resource-constrained, market-facing sectors within Zambia and the wider region, including private healthcare providers and financial-services institutions, that are adopting AI tools under comparable regulatory ambiguity. Public higher education institutions, which typically operate under closer ministry oversight and more developed quality-assurance infrastructure, may face a narrower version of the same gap, and testing whether the pattern documented here weakens, holds, or intensifies under public-sector governance conditions is a natural extension of this paper's approach. These extensions are offered as plausible directions for comparative research rather than as claims this single-country, single-subsector study can itself support.

These recommendations are deliberately sequenced rather than presented as a simultaneous checklist. Literacy development and governance-framework development are foundational and achievable at the institutional level without waiting on external funding or regulatory change, and should accordingly come first. Regulatory and policy responses, which require coordination beyond any single institution's control, appropriately follow. Attempting sector-wide regulatory mandate before institutions have begun building the underlying literacy and governance capacity risks producing compliance documents that exist on paper without the leadership understanding needed to operate them, echoing the same literacy-governance disconnection this paper's findings describe at the individual level.

This study's findings should be read alongside several limitations. The qualitative interviews, while purposively varied by institution type, size, and ownership, were collected at five of the sector's twelve Tier 1 universities, and the specific institutional accounts reported here should not be assumed to generalise without qualification to the remaining seven. The cross-sectional design captures a single point in a rapidly evolving technological and regulatory landscape, and both individual literacy and institutional response are plausibly moving targets rather than fixed states. Self-report measures of literacy and institutional practice, while corroborated here by consistent and self-implicating qualitative accounts that a social-desirability bias would not straightforwardly predict, remain subject to the general limitations of self-assessment. Future research would usefully track literacy and institutional governance longitudinally as capacity-building interventions are introduced, and would extend the qualitative inquiry to the remaining Tier 1 institutions to test whether the governance patterns described here hold sector-wide.

Ethical Considerations

This study received ethical clearance from the Africa Research University Research Ethics Committee (Certificate Reference ARU/2026/05/005; approved 29 May 2025). All data collection took place within the approval period. Prior to participation, all respondents and interview participants received a written participant information sheet describing the study's purpose, the nature and voluntary basis of participation, how data would be used, and the right to withdraw at any time without consequence; written informed consent was obtained before any data collection began, and no coercion or inducement was used in recruitment. Institutional access was obtained through formal written requests to each participating university. Questionnaire responses were collected anonymously, with no identifying information recorded. Interview participants were assigned pseudonymous codes (P1 to P12) and are identified throughout only by role category; no institutional affiliation is attached to any individual participant account, reflecting the small number of senior office-holders per

institution and the consequent risk of indirect identification. Audio recordings and transcripts were stored in password-protected files on an encrypted device, accessible only to the research team, with signed consent documentation stored separately from research data. All data will be retained securely for five years following publication, in accordance with Africa Research University's data management policy, and thereafter securely destroyed. Interview questions elicited professional governance perspectives rather than personal information, and no participant was exposed to foreseeable physical, psychological, professional, or reputational risk through participation.

ACKNOWLEDGEMENTS

The authors thank the strategic leaders and institutional staff across Zambia's Tier 1 private universities who gave their time to complete the questionnaire and participate in interviews for this study, and Africa Research University for institutional support during the research.

Competing Interests

The authors declare no competing interests. The corresponding author's institutional affiliation, as Principal and Registrar at one of the twelve Tier 1 institutions constituting the study's target population, is disclosed as a matter of transparency; this insider position is addressed directly as part of the study's methodology and reflexivity practice.

Author Contributions

I. Kumwenda: conceptualisation, methodology, data collection, formal analysis, writing – original draft. J. Likando: supervision, methodology review, writing – review and editing. M. Malama: data curation, validation, writing – review and editing. H. Siame: investigation, validation, writing – review and editing. D. Mainde: resources, writing – review and editing. All authors approved the final version of the manuscript and agree to be accountable for its content.

Funding

This research received no external or institutional grant funding. Data collection costs were met by the corresponding author.

Data Availability

The questionnaire and interview data supporting this study are not publicly available due to the confidentiality undertakings given to participants and the institutions involved but are available from the corresponding author on reasonable request, subject to ethical approval and institutional consent.

Disclaimer

The views expressed in this article are those of the authors and do not necessarily represent the official position of Africa Research University or any of the institutions whose staff participated in this study.

REFERENCES

1. Anderson, F.V. (2025) 'Developing AI leadership capacity in Africa: exploring the role of education, training, and mentorship programs', IRPJ (Intergovernmental Research and Policy Journal). Available at: <https://irpj.eu>
2. Bass, B.M. and Avolio, B.J. (1994) *Improving Organizational Effectiveness Through Transformational Leadership*. Thousand Oaks, CA: Sage Publications.
3. Braun, V. and Clarke, V. (2006) 'Using thematic analysis in psychology', *Qualitative Research in Psychology*, 3(2), pp. 77–101.
4. Braun, V. and Clarke, V. (2022) *Thematic Analysis: A Practical Guide*. London: Sage Publications.
5. Bryman, A. (2016) *Social Research Methods*. 5th edn. Oxford: Oxford University Press.

6. Buolamwini, J. and Gebru, T. (2018) 'Gender shades: intersectional accuracy disparities in commercial gender classification', *Proceedings of Machine Learning Research*, 81, pp. 77–91.
7. Burns, J.M. (1978) *Leadership*. New York: Harper & Row.
8. Cohen, L., Manion, L. and Morrison, K. (2018) *Research Methods in Education*. 8th edn. London: Routledge.
9. Creswell, J.W. and Plano Clark, V.L. (2023) *Designing and Conducting Mixed Methods Research*. 4th edn. Thousand Oaks, CA: Sage Publications.
10. Davis, F.D. (1989) 'Perceived usefulness, perceived ease of use, and user acceptance of information technology', *MIS Quarterly*, 13(3), pp. 319–340.
11. DiMaggio, P.J. and Powell, W.W. (1983) 'The iron cage revisited: institutional isomorphism and collective rationality in organizational fields', *American Sociological Review*, 48(2), pp. 147–160.
12. Dube, B., Nkomo, D. and Apadile-Thokweng, M. (2024) 'Pragmatism: an essential philosophy for mixed methods research in education', *International Journal of Research and Innovation in Social Science*, 8(3), pp. 1001–1010.
13. Dusengumuremyi, F., Chukwudi John, H., Ndahimana, S., Ndahayo, C. and Sangwa, S. (2026) 'From innovation signaling to governance readiness: generative artificial intelligence integration, policies, opportunities, and challenges in two Rwandan universities', *Open Journal of AI, Ethics and Society*, 2(2).
14. Fiedler, F.E. (1967) *A Theory of Leadership Effectiveness*. New York: McGraw-Hill.
15. Hersey, P. and Blanchard, K.H. (1982) *Management of Organizational Behavior: Utilizing Human Resources*. 4th edn. Englewood Cliffs, NJ: Prentice-Hall.
16. Higher Education Authority of Zambia (2024) *The State of Higher Education in Zambia 2023: Quality Assurance in Higher Education*. Lusaka: Higher Education Authority.
17. Higher Education Authority of Zambia (2026a) *Gazette Notice No. 665 of 2026: Classification of Universities*. Government Gazette of the Republic of Zambia, Vol. LXII, No. 46, 17 April 2026. Lusaka: Government Printer.
18. Jacobs-Basadien, M. and Pather, S. (2026) 'Assessing AI policy and guideline progress within the South African public higher education system: insights from documented policies', *Universal Access in the Information Society*, 25(2), article 71. Available at: <https://doi.org/10.1007/s10209-026-01340-9>.
19. Kvale, S. and Brinkmann, S. (2015) *InterViews: Learning the Craft of Qualitative Research Interviewing*. 3rd edn. Thousand Oaks, CA: Sage Publications.
20. Long, D. and Magerko, B. (2020) 'What is AI literacy? Competencies and design considerations', in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery.
21. Memarian, B. and Doleck, T. (2023) 'Fairness, accountability, transparency, and ethics (FATE) in artificial intelligence (AI) and higher education: a systematic review', *Computers and Education: Artificial Intelligence*, 5.
22. Meyer, J.W. and Rowan, B. (1977) 'Institutionalized organizations: formal structure as myth and ceremony', *American Journal of Sociology*, 83(2), pp. 340–363.
23. Miao, F., Holmes, W., Huang, R. and Zhang, H. (2021) *AI and Education: Guidance for Policy-Makers*. Paris: UNESCO.
24. Miles, M.B., Huberman, A.M. and Saldaña, J. (2020) *Qualitative Data Analysis: A Methods Sourcebook*. 4th edn. Thousand Oaks, CA: Sage Publications.
25. Nunnally, J.C. (1978) *Psychometric Theory*. 2nd edn. New York: McGraw-Hill.
26. O'Neil, C. (2016) *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown Publishers.
27. Pasipamire, N. and Muroyiwa, A. (2024) 'Navigating algorithm bias in AI: ensuring fairness and trust in Africa', *Frontiers in Research Metrics and Analytics*, 9, 1486600. Available at: <https://doi.org/10.3389/firma.2024.1486600>.
28. Pretorius, L. (2024) 'Demystifying research paradigms: navigating ontology, epistemology, and axiology in research', *The Qualitative Report*, 29(10), pp. 2698–2715.
29. Sangwa, S., Ngobi, D., Ekosse, E. and Mutabazi, P. (2026) 'AI governance in African higher education: status, challenges, and a future-proof policy framework', *Artificial Intelligence and Education*, 1(1), Article 2054. Available at: <https://doi.org/10.62617/AIE2054>.

30. Saunders, M., Lewis, P. and Thornhill, A. (2019) *Research Methods for Business Students*. 8th edn. Harlow: Pearson.
31. Venkatesh, V. and Bala, H. (2008) 'Technology Acceptance Model 3 and a research agenda on interventions', *Decision Sciences*, 39(2), pp. 273–315.
32. Venkatesh, V., Morris, M.G., Davis, G.B. and Davis, F.D. (2003) 'User acceptance of information technology: toward a unified view', *MIS Quarterly*, 27(3), pp. 425–478.
33. Weiner, B.J. (2009) 'A theory of organizational readiness for change', *Implementation Science*, 4, article 67. Available at: <https://doi.org/10.1186/1748-5908-4-67>.
34. Zambia, Ministry of Technology and Science (2025) *National Artificial Intelligence Strategy 2025–2027*. Lusaka: Government of the Republic of Zambia.