SAMYRAD 2026

SMART ALGORITHMS & MODELS: YOUR REACH TO ADVANCED DIALOGUES · 5–6 OCTOBER 2026 · SEVILLE, SPAIN

SAMYRAD OPEN-ACCESS PROCEEDINGS | PUBLISHER: INTELLIGENT NUDGE

PAPER • OPEN ACCESS

From Static LMS to Autonomous Agents: Reimagining Moodle-Based Open and Distance e-Learning Delivery for the Agentic AI Era in Kenya

Sadati Okeyo[1], Joel Nakitare[2]

(email. sokeyo@umma.ac.ke), (email. jnakitare@rongovarsity.ac.ke)

[1]Systems Librarian Umma University Kajiado, Kenya

[2]University Librarian Rongo University Rongo, Kenya

Article DOI: . https://doi.org/10.68286/samyrad2026.348

ABSTRACT

Artificial intelligence (AI) in education has progressed from generative content tools toward agentic systems capable of autonomous planning, decision-making and multi-step task execution. Learning Management Systems (LMS) such as Moodle have underpinned Open and Distance e-Learning (ODeL) delivery across Kenyan higher education institutions, yet their architecture remains largely static, offering limited capacity for autonomous, adaptive learner support. This study addressed a critical gap: the absence of empirical evidence on institutional readiness for agentic AI integration within Moodle-based ODeL ecosystems in Kenya. A descriptive cross-sectional survey was administered online to respondents drawn from institutional leadership, ICT/eLearning directors, ODeL coordinators and registered ODeL students at Kenyan higher education institutions offering ODeL. Seventeen (N = 17) respondents took part, with question-level completion ranging from 8 to 17 owing to conditional branching and voluntary attrition. Findings show that 88.9% of respondents whose institutions had customised Moodle reported moderate-to-full customisation, and 87.5% reported that their institution had piloted or implemented at least one adaptive or automated function, most commonly automated plagiarism detection (100%) and adaptive learning paths (71.4%); because current implementations cluster around integrity and pathway automation rather than autonomous, goal-directed agents, these findings are best read as evidence of preparedness for a prospective transition to agentic AI rather than of its current widespread deployment. A Fisher's exact test found no statistically significant association between perceived ICT budget adequacy and adoption of adaptive functionality ($p = 1.00$), and a Welch's independent samples-t-test found no statistically significant difference in perceived institutional readiness between staff ($M = 3.25$, $SD = 1.71$) and students ($M = 4.50$, $SD = 0.58$), $t(3.68) = -1.39$, $p = .244$, although the small sample severely limits statistical power. Infrastructural constraints, weak institutional AI policy, and equity and data-governance concerns emerged as the principal qualitative barriers. The study proposes a four-domain conceptual readiness framework, specified a priori from the technology-adoption and governance literature (infrastructure, policy and governance, human capability, and equity), to guide Kenyan ODeL institutions in transitioning from static LMS infrastructure toward agentic-AI-driven delivery, in alignment with Sustainable Development Goal 4. Given the modest and private-institution-skewed sample, the findings are best read as an exploratory baseline rather than a population-representative estimate, and the paper closes with a call for a larger, multi-institution replication.

KEYWORDS agentic AI, Moodle , ODeL, learning management systems, Kenya, higher education, institutional readiness

How to cite: Okeyo, S., & Nakitare, J. (2026). From static LMS to autonomous agents: Reimagining Moodle-based open and distance e-learning delivery for the agentic AI era in Kenya. SAMYRAD 2026 Open-Access Proceedings. Intelligent Nudge. https://doi.org/10.68286/samyrad2026.348

Introduction

The global higher-education technology landscape has shifted decisively from generative AI tools that respond to prompts toward agentic AI systems that plan, act and adapt with minimal human supervision [1], [4]. Agentic AI is distinguished from earlier intelligent tutoring systems and conversational assistants by its capacity to pursue multi-step, long-horizon goals, coordinate tools and agents, and adjust its own plans as circumstances change [2], [4], [5]. In higher education, this is beginning to manifest as autonomous learning companions, adaptive assessment agents and proactive student-support agents that act on institutional data rather than merely displaying it [6].

Kenyan higher education institutions have widely adopted Moodle as the backbone of Open and Distance e-Learning (ODeL) delivery, a trend accelerated by the COVID-19 pandemic and by continuing demand that outstrips the capacity of conventional face-to-face provision [7], [8], [9]. However, Moodle deployments in the region have historically served as repositories for content delivery and administrative workflow rather than as platforms for autonomous, adaptive learner support [7]. Whether Kenyan ODeL institutions currently possess the technical, human, policy and equity foundations required to responsibly adopt agentic AI within Moodle remains an open empirical question.

This study addressed that gap directly. It asked: (i) to what extent have Kenyan ODeL institutions customised Moodle and piloted adaptive or automated functionality; (ii) how adequate do stakeholders judge the ICT budget, infrastructure, policy and governance foundations for agentic AI; (iii) do perceptions of readiness differ between institutional staff and students; and (iv) what do stakeholders identify as the principal barriers to, and enablers of, a transition toward agentic-AI-driven Moodle delivery. The paper contributes an empirically grounded, exploratory baseline and a four-domain conceptual readiness framework intended to guide systems librarians, ICT directors and ODeL administrators as they plan this transition, in a manner consistent with Sustainable Development Goal 4 (quality education).

The contribution of this paper is threefold. It provides, to the authors' knowledge, an early empirically grounded readiness assessment of agentic AI within Moodle-based ODeL delivery specific to the Kenyan context, at a moment when agentic functionality is moving rapidly from pilot to mainstream deployment internationally [5], [6]. It demonstrates, transparently and with appropriate statistical caution, what a small-sample institutional survey can and cannot support, a methodological contribution relevant to systems librarians and ODeL administrators who often work with similarly constrained data-collection capacity. Finally, it offers a four-domain conceptual framework, specified a priori from the Technology, Organisation and Environment (TOE) tradition [14], [15] rather than induced from the survey data, that later, larger-scale studies can operationalise and test.

Related work

Agentic AI in Education

Agentic AI is consistently defined in the recent literature as AI that pursues complex, often long-horizon goals with minimal human intervention, planning, using tools, maintaining memory and self-adjusting to changing environments [1], [2]. This marks a qualitative shift from earlier intelligent tutoring systems, which operated within a fixed curriculum, and from generative AI assistants, which remain fundamentally reactive to user prompts [4]. Surveys of the field classify emerging applications into intelligent tutoring, adaptive assessment, autonomous learning companions and multi-agent orchestration layers that sit across an institution's existing systems [2], [3]. Commentators note that the shift from copilot to autopilot modes of operation raises new questions of trust, transparency and governance that are only beginning to be addressed in policy [1], [3], [5], and that adoption in the sector is accelerating rapidly even as organisational trust in autonomous systems lags behind investment [5], [6].

Moodle and ODeL in Kenya

Kenyan and regional studies confirm that Moodle is now a common LMS across the higher-education sector, frequently supplemented by institution-specific customisation [7]. Institutional-factors research on Kenyan e-learning adoption identifies organisational leadership support, staff readiness and resourcing as significant predictors of successful technology transitions [8], while national e-readiness surveys commissioned through the Kenya Education Network document persistent gaps in device-to-student ratios and ICT human capacity across the university sector [9]. Implementation-focused work further reports that device and connectivity access remain uneven, particularly at public universities, and recommends participatory planning involving leadership, staff and students, together with explicit institutional ODeL policy, as preconditions for successful technology transitions [10], a recommendation this study revisits in the specific context of agentic AI.

Data Governance and Equity

Agentic AI systems that act on learner data raise data protection obligations that did not attach to static LMS use in the same way. Kenya's Data Protection Act, 2019 requires lawful, transparent and accountable processing of personal data and establishes the Office of the Data Protection Commissioner (ODPC) as the enforcing authority [11]; the ODPC's sector guidance note for education explicitly extends these obligations to institutions offering remote e-learning [12]. Separately, equity scholarship on African distance education consistently documents an urban and rural digital divide in device ownership, bandwidth and consequently in the ability to benefit from technology-enhanced learning, with rural and historically disadvantaged students most exposed [13]. Any readiness assessment for agentic AI in an African ODeL context must therefore treat data governance and equity as first-order dimensions rather than afterthoughts.

Theoretical Grounding: Institutional Readiness Models

Institutional technology adoption is commonly theorised through the Technology, Organisation and Environment (TOE) framework, which holds that adoption outcomes are jointly shaped by technology characteristics (complexity, compatibility, relative advantage), organisational characteristics (leadership support, financial resources, internal readiness) and environmental characteristics (policy pressure, regulatory context, competitor behaviour) [14]. Applying TOE specifically to generative AI adoption in higher education, Kalmus and Nikiforova found that environmental and organisational constraints, including regulatory ambiguity and unclear institutional policy, shaped educator resistance even where individual attitudes toward the technology were otherwise favourable [15], echoing this study's own finding that policy and governance items lagged infrastructure items despite generally favourable individual attitudes. This study's four thematic blocks (infrastructure, budget/organisation, policy/environment, and equity) were specified a priori from this literature, before data collection, and are used to structure the proposed conceptual framework as a literature-derived contribution rather than a pattern induced from the survey data.

AI Governance and Accountability in Higher Education

Agentic systems complicate conventional accountability structures because a single student-facing outcome may result from a multi-step, tool-using pipeline rather than a single identifiable human decision [3], [16]. Emerging governance scholarship addresses this by proposing that a named human accountability owner be assigned to each deployed agentic workflow and that agents be granted scoped, time-limited permissions rather than static, standing access to personal data, extending established privacy-engineering principles to autonomous, tool-using systems [16]. These developments reinforce why this study treats a formal AI-in-teaching policy and an explicit accountability framework as first-order readiness indicators rather than administrative afterthoughts.

Methodology

Research Design

The study adopted a descriptive, cross-sectional survey design, appropriate for establishing a baseline profile of institutional readiness across multiple stakeholder perspectives at a single point in time. The design combined closed-ended items (single-choice, multiple-choice and five-point Likert matrices) with open-ended items to allow both quantifiable comparison and qualitative elaboration of barriers and support needs.

Population, Sampling and Participants

The target population comprised stakeholders at Kenyan higher education institutions that had formally embraced ODeL delivery, spanning four strata: institutional leadership (Deans/Directors of ODeL), ICT/eLearning directors, ODeL coordinators, and registered ODeL students. Purposive and snowball sampling within these strata were used to reach respondents through institutional ODeL and ICT departments. The instrument was administered online via SurveyMonkey between 20 June and 7 July 2026 and yielded seventeen respondents (N = 17): 11.8% from public institutions and 88.2% from private institutions, distributed across the four strata as registered ODeL students (53.3%), ICT/eLearning directors (20.0%), ODeL coordinators (20.0%) and institutional leadership (6.7%) among the fifteen respondents who disclosed a category. The pronounced private-institution skew is a material feature of the sample, not merely a demographic footnote, and is used in the Results Section (Budget, Infrastructure, Policy and Equity Perceptions) to explain an otherwise counter-intuitive result. The survey did not capture a unique institutional identifier, so N = 17 is reported as individual respondents rather than as a count of distinct institutions, and demographic detail beyond institution type and respondent category (e.g., years of experience, age) was not collected; both points are addressed in the Limitation Section.

Instrument

The questionnaire comprised eight thematic blocks: respondent and institution profile; Moodle customisation; adaptive/automated functionality already piloted; ICT budget and its adequacy; four five-point Likert agreement matrices covering (i) overall readiness, (ii) infrastructural barriers, (iii) policy and data-governance maturity, and (iv) equity; and three open-ended questions on barriers, support needed, and additional comments. The four Likert blocks correspond directly to the four framework domains specified a priori in the Related Work Section (Theoretical Grounding: Institutional Readiness Models) and the framework is later discussed as a literature-derived structure that the instrument operationalises, not as a finding the analysis selected among alternatives. Likert items were anchored 1 = Strongly disagree to 5 = Strongly agree. Skip logic meant later blocks were shown only to respondents who indicated relevant prior experience (e.g., the functionality-detail block was shown only to those who had piloted adaptive features), which explains the declining n across blocks (17 down to as few as 6 or 8) reported in the Results Section.

Data Analysis

Closed-ended responses were analysed descriptively (frequencies, percentages, means and standard deviations) using the response summaries generated in SurveyMonkey. Two inferential tests were run on the fully-completed Likert and categorical items (n = 8): a Welch's independent-samples t-test comparing perceived overall institutional readiness between staff (institutional leadership, ICT directors and ODeL coordinators pooled) and students, and a Fisher's exact test examining the association between perceived ICT budget adequacy (collapsed into adequate/neutral versus inadequate/very inadequate) and whether the institution had piloted adaptive functionality. Welch's t-test and Fisher's

exact test were selected in preference to the pooled-variance t-test and the standard Pearson chi-square test because both are more robust to the unequal variances and small expected cell counts (< 5) that characterise a sample of this size; a Mann-Whitney U test and an uncorrected chi-square are reported alongside as robustness checks. Degrees of freedom for the Welch's t-test were computed using the Welch-Satterthwaite equation rather than assumed from sample size, since the two are not interchangeable under unequal variances. All descriptive standard deviations are reported using the sample (n-1) formula for consistency with the inferential statistics. Open ended responses were analysed thematically by inductive coding of recurring terms. All quantitative analysis was performed in Python 3 using SciPy 1.x.

Ethical Considerations

Content validity was addressed by deriving each thematic block directly from constructs established a priori in the adoption and equity literature reviewed in the Related Work Section, namely TOE-aligned infrastructure, organisational and environmental items [14], [15], data-governance items aligned with the Kenya Data Protection Act, 2019 [11], [12], and equity items drawn from the African digital divide literature [13]. Given the exploratory nature and small sample of this study, formal internal-consistency statistics (e.g., Cronbach's alpha) were not computed for the Likert matrices, since such coefficients are unstable and can be misleading at n = 8; this is flagged explicitly as a target for the larger replication study recommended in the Future Research Directions Section. Face validity was supported by piloting the questionnaire informally with two ODeL practitioners prior to distribution, who confirmed that item wording was unambiguous in the Kenyan institutional context.

Validity and Reliability

Participation was voluntary and anonymous. The instrument's introductory page informed respondents that the survey was part of an academic research study, that responses were confidential, and that completion was expected to take twelve to fifteen minutes; no personally identifying information (name, email address, or institution name) was collected. Because the Content validity was addressed by deriving each thematic block directly from constructs established a priori in the adoption and equity literature reviewed in the Related Work, namely TOE-aligned infrastructure, organisational and environmental items [14], [15], data-governance items aligned with the Kenya Data Protection Act, 2019 [11], [12], and equity items drawn from the African digital divide literature [13]. Given the exploratory nature and small sample of this study, formal internal-consistency statistics (e.g., Cronbach's alpha) were not computed for the Likert matrices, since such coefficients are unstable and can be misleading at n = 8; this is flagged explicitly as a target for the larger replication study as recommended in 'Proposed conceptual readiness framework' section. Face validity was supported by piloting the questionnaire informally with two ODeL practitioners prior to distribution, who confirmed that item wording was unambiguous in the Kenyan institutional context did not record institutional identity, individual institutions cannot be singled out from the published results, which further protects respondent confidentiality but also constrains the analysis, as discussed in the Limitation Section. Respondents could exit the survey at any point, which accounts for the partial-completion pattern evident in the declining question-level n reported in the Results Section.

Results and findings

Respondent and Institutional Profile

Table 1 summarises the respondent and institutional profile; while the report of the completion count (n) for every item discussed below is presented later so that percentages computed over different respondent subsets remain traceable. The sample was dominated by private-institution respondents (88.2%) and by registered ODeL students (53.3%), a skew used to interpret an otherwise counter-intuitive result and flagged as a limitation of the study.

Table 1. Respondent and Institutional Profile

Element	Example value	n	%
Institution type (n=17)	Public	2	11.8
	Private	15	88.2
Respondent category (n=15) Leadership 1 6.7	Leadership	1	6.7
	ICT/eLearn. Dir.	3	20.0
	ODeL Coord.	3	20.0
	ODeL Student	8	53.3

traceable. The sample was dominated by private-institution respondents (88.2%) and by registered ODeL students (53.3%), a skew used in the Results Section (Budget, Infrastructure, Policy and Equity Perceptions) to interpret an otherwise counter-intuitive result and flagged as a limitation.

Moodle Use and Customisation

Of the ten respondents who answered, 90.0% confirmed Moodle as their institution's primary LMS for ODeL delivery. Among the nine who detailed the extent of customisation, 55.6% reported extensive customisation integrated with other systems, 22.2% moderate customisation (plugins added), 11.1% full customisation with AI-enhanced plugins, and 11.1% no customisation beyond default installation. Among the ten who itemised specific customisations (multiple selection permitted), additional plugins (80.0%), custom theming (70.0%) and API integrations (70.0%) were the most common, followed by analytics dashboards (60.0%), third-party integrations such as payment or SMS gateways (50.0%) and mobile app integration (40.0%); only 10.0% reported no customisation at all.

Adaptive and Automated Functionality

Among the eight respondents who answered, 87.5% reported that their institution had piloted or implemented at least one adaptive or automated learning function within Moodle, against 12.5% who had not. Table 2 shows the specific functionalities reported by the seven respondents who answered "yes" and went on to itemise which functionalities were in use; the deployment-stage question that follows was instead shown to all eight respondents who answered the piloting question, including the one respondent who had not piloted any functionality and who selected "discontinued use", which accounts for the differing n between these two items. Automated plagiarism detection was universal among adopters (100%), followed by adaptive learning paths (71.4%) and automated grading (57.1%); chatbot/virtual-assistant support and AI-generated content were each reported by only 14.3%, suggesting that current "adaptive" practice in Kenyan Moodle deployments still clusters around integrity and pathway tools rather than conversational or generative agentic functions, and that the readiness this paper assesses is preparedness for a prospective transition rather than evidence of agentic AI already being widely deployed. Of the eight respondents who indicated a deployment stage, 50.0% reported full institutional deployment, 25.0% partial rollout, 12.5% pilot/testing and 12.5% discontinued use (Table 2).

Table 2. Adaptive/Automated Functionalities Piloted (n=7, multi-select)

Functionality	%
Automated plagiarism detection	100
Adaptive learning paths	71.4
Automated grading	57.1
Predictive analytics for at-risk students	42.9
Chatbot / virtual assistant support	14.3
AI-generated content / quizzes	14.3

Budget, Infrastructure, Policy and Equity Perceptions

Budget visibility was low: of eight respondents, 62.5% (five) did not know what share of the institutional budget was allocated to ICT, and 75.0% (six) did not know whether a dedicated LMS/AI innovation budget line existed; the remaining 25.0% (two respondents) confirmed that a dedicated line did exist, and none reported affirmatively that no such line existed. Perceived budget adequacy for agentic AI specifically was mixed: 37.5% rated it neutral, 25.0% adequate, 25.0% very inadequate, and 12.5% inadequate, with none rating it very adequate.

Table 3 reports mean agreement (on a 1 to 5 scale, sample standard deviation) across the four Likert matrices (n = 8 throughout). Readiness-related items clustered in the neutral-to-agree range (M = 3.50 to 3.88), as did policy and governance items, with self-reported compliance with the Data Protection Act, 2019 rated highest overall (M = 4.13). Notably, the infrastructural-barrier items were rated closer to the disagree end of the scale (M = 2.25 to 3.00), indicating that respondents in this sample did not, on balance, regard connectivity, device shortages, hosting capacity, power outages or technical-support gaps as acute; this is a more optimistic picture than the barriers literature on Kenyan e-learning generally reports [8], [9], [10], and is most plausibly explained by the sample's 88.2% private institution skew (Table 1), since private institutions in Kenya are generally better resourced than the public universities on which that literature disproportionately reports, rather than by sample size alone. Equity items were more concerning: 75.0% agreed or strongly agreed that rural or remote students face disadvantages in accessing agentic AI features (M = 3.75), the highest mean of any equity item.

Table 3. Mean agreement by domain (1=strongly disagree, 5=strongly agree, n=8, sample SD)

Item	M	SD
Institution has technical infrastructure for agentic AI	3.63	0.75
Staff have adequate AI-management skills	3.75	0.89
Students prepared for autonomous AI tools	3.75	0.89
Institution has clear AI strategic plan	3.5	0.93
Leadership visibly committed to AI adoption	3.63	1.19
Confident institution can manage data-privacy risk	3.50	0.76
Overall: institution ready for agentic-AI Moodle	3.88	1.36
Unreliable internet hampers LMS use	2.75	1.17
Insufficient computing devices	2.50	1.20
Server/hosting capacity inadequate	2.25	1.28
Frequent power outages disrupt learning	3.00	0.93
Lack of dedicated technical support staff	2.25	1.04
Institution has formal AI-in-teaching policy	3.75	1.28
Clear data-governance guidelines for AI data	3.25	1.04
Complies with Kenya Data Protection Act (2019)	4.13	0.83
Clear accountability framework for AI decisions	3.63	0.92
Equal device/internet access for AI-enhanced learning	3.25	1.04
Students with disabilities adequately supported	3.00	1.30
Rural/remote students disadvantaged in AI access	3.75	0.46

Staff vs. Student Perceived Readiness

Restricting to the eight fully-completed responses, staff (institutional leadership, ICT directors and ODeL coordinators, n = 4) reported an overall-readiness mean of 3.25 (SD = 1.71; individual scores 1, 3, 4, 5), while students (n = 4) reported a higher mean of 4.50 (SD = 0.58; scores 4, 4, 5, 5). A Welch's independent-samples t-test, with degrees of freedom computed via the Welch-Satterthwaite equation rather than assumed from group size, found this difference

was not statistically significant, t(3.68) = -1.39, p = .244; a Mann-Whitney U test corroborated this, U = 4.0, p = .288. The mean difference was 1.25 points on the five-point scale, 95% CI [-1.34, 3.84], with a large point-estimate effect size (Hedges' g = 0.85), but the interval spans zero and covers both a negligible and a very large true effect, the expected signature of a real difference tested with too few observations to pin down, not evidence that no difference exists. The direction of the gap (students were somewhat more optimistic than staff) is the reverse of what is often assumed in the technology-adoption literature and of what the original study proposal anticipated, but with n = 4 per group and fewer than four whole degrees of freedom the test is severely underpowered and this result should be treated as inconclusive rather than as evidence of no true difference.

Budget Adequacy and Adoption of Adaptive Functionality

A 2×2 contingency table (n = 8) cross-tabulating collapsed budget adequacy (adequate/neutral vs. inadequate/very inadequate) against whether the institution had piloted adaptive functionality (yes/no) is shown in Table 4. A Fisher's exact test (preferred here because two of the four expected cell counts fell below 5) found no statistically significant association, p = 1.00, odds ratio = 0.00. An uncorrected Pearson chi-square test is reported for reference only, χ^2 (1, n = 8) = 0.69, p = .408; given the small expected counts this chi-square result should not be relied upon. In short, this sample provides no evidence that perceived budget adequacy is associated with whether adaptive functionality has been piloted; this null result is plausibly attributable to the very small sample rather than to a genuine absence of relationship, and it reverses the significant association hypothesised in the original study proposal.

Table 4. Budget Adequacy × Adaptive-Functionality Adoption (n=8)

Budget adequacy (collapsed)	Piloted: Yes	Piloted: No
Adequate / Neutral	4	1
Inadequate / Very inadequate	3	0

Qualitative Barriers and Support Needs

Six respondents identified the single greatest barrier to agentic AI adoption. Data privacy, security and regulatory-compliance concerns were the most frequently cited theme, followed by inadequate staff skills and technology apprehension, gaps in leadership attention and resourcing, general lack of awareness, and technical/systems failures. On support needed (n = 6), financial investment was named in almost every response, typically paired with technical capacity-building or training, and with calls for clearer AI governance policy. Additional open comments (n = 5) emphasised that adoption should be approached strategically and responsibly, with attention to staff readiness and data security, and that sensitisation ahead of adoption was necessary; one respondent offered no further comment and one described the topic positively without elaboration (Table 5).

Table 5. Sample Size (n) by Survey Item

Survey block	Piloted: No
Moodle as primary LMS	10
Extent of customisation	9
Specific customisations (multi-select)	10
Piloted adaptive functionality (yes/no)	8
Specific functionalities piloted (multi-select)	7
Deployment stage	8
ICT budget share / dedicated line / adequacy	8 each
Four Likert matrices (19 items)	8 each
Open-ended barriers / support / comments	6 / 6 / 5

Discussion

Three findings stand out. First, Kenyan ODeL institutions in this sample have already moved well beyond a default installation of Moodle: extensive customisation, API integrations and analytics dashboards are common, and a majority report having piloted at least one adaptive or automated function, most often plagiarism detection and adaptive pathways rather than conversational or generative agents. This suggests the technical substrate for agentic AI (APIs, integrations, data pipelines) is more developed in this sample than the static-LMS framing in prior literature might imply, though functionality remains concentrated in integrity and routing tools rather than autonomous, goal-directed agents as defined in the agentic AI literature [1], [2], [4]; the study therefore measures readiness for a prospective transition, not the extent of current agentic deployment.

Second, neither of the two inferential hypotheses tested here (a staff-student readiness gap and a budget-adoption association) was supported at conventional significance levels. This null result is itself informative: it cautions against assuming, without institution-level data, that budget alone drives adoption, or that frontline staff are systematically more sceptical than students. It also underscores a structural limitation of this study, discussed in the Limitation Section, rather than a substantive conclusion that can be generalised beyond this sample.

Third, the qualitative and equity data together paint a more cautious picture than the readiness means alone. Data governance and compliance concerns dominate the open-ended barrier responses even though self-reported Data Protection Act compliance was the single highest-rated item on the closed-ended matrices (M = 4.13); this is a tension that suggests respondents may distinguish formal legal compliance from confidence in operational data governance for autonomous systems specifically.

Similarly, 75% agreement that rural and remote students face disadvantages in accessing agentic AI features echoes the wider African ODeL equity literature [13] and signals that

agentic AI, if adopted without deliberate equity safeguards, risks widening rather than narrowing existing access gaps.

Read through a TOE lens, the findings suggest an uneven profile across the three dimensions [14], [15]. The technology dimension appears comparatively mature, evidenced by widespread API integration, analytics dashboards and third-party integrations. The organisational dimension is mixed: leadership commitment and staff skill were both rated moderately, yet budget visibility was strikingly low, with roughly two-thirds of respondents unable to say what share of the institutional budget supports ICT. The environmental dimension, comprising policy and regulatory context, was the most consistently flagged as unresolved in the open-ended responses, even though formal legal compliance with the Data Protection Act was rated highly on the closed-ended items. This pattern is consistent with Kalmus and Nikiforova's TOE-informed finding that environmental and governance constraints, rather than technology or organisational readiness alone, most often shape resistance to AI adoption in higher education [15].

For practice, the implication is that institutions preparing for agentic AI in Moodle should not treat budget allocation as a proxy for readiness, given the null association observed in the Results Section (Budget Adequacy and Adoption of Adaptive Functionality) and the low-budget visibility reported by respondents; a dedicated, transparent LMS/AI budget line, tracked and communicated to staff, may do more to signal institutional readiness than the raw size of the ICT budget. The emerging practice of naming a human accountability owner for each deployed agentic workflow, and granting agents scoped rather than static permissions to student data [16], offers a concrete governance mechanism that Kenyan ODeL institutions could adapt directly, and would address the accountability gap that respondents flagged even where formal legal compliance was rated highly.

Limitations

Sample size and power: N = 17 overall, with usable n for inferential tests as low as 4 per group, severely limits statistical power; the null results reported in the Results Sections (Staff vs. Student Perceived Readiness and Budget Adequacy and Adoption of Adaptive Functionality) should be read as inconclusive, not as evidence of absence. Institutional identification: the instrument did not capture a unique institution identifier, so the sample cannot be characterised with confidence as spanning a known number of distinct institutions, and clustering of respondents within the same institution cannot be ruled out. Sample composition: the sample is skewed toward private institutions (88.2%) and registered students (53.3%), while institutional leadership was represented by a single respondent; the Result Section (Budget, Infrastructure, Policy and Equity Perceptions) shows this skew, rather than sample size alone, most plausibly explains the sample's unexpectedly benign infrastructural-barrier ratings, and it limits how far any finding can be generalised to the public universities that dominate national ODeL provision. Attrition and skip logic: question-level n fell from 17 to as low as 6 or 8 for later blocks introducing potential non-response bias toward more engaged or more technically confident respondents. Self-report: all measures, including infrastructural barrier items, rely on respondent perception rather than independently verified institutional data (e.g., audited ICT budgets, measured connectivity). Respondent detail: beyond institution type and respondent category, the instrument did not collect further demographic detail (e.g., years of experience, age, gender), limiting the granularity with which sub-groups can be described or compared.

Proposed Conceptual Readiness Framework

Status of the Framework and Its Relation to Prior Models

The four domains presented in this section were specified a priori from the technology-adoption and governance literature reviewed (principally TOE [14], [15] and the AI-governance literature [3], [16]), and structured both the survey instrument and the discussion. They are a literature-derived conceptual contribution, not an inductive finding the survey data selected over alternative groupings; the close correspondence between the domains and the instrument's Likert blocks reflects that the instrument operationalises the framework, not the reverse. One reasonable objection is that TOE's environmental dimension could already accommodate regulatory and equity pressures, making a separate equity domain redundant with TOE. We retain equity as an explicit, cross-cutting domain because, in the pre-scaling audit use this framework targets equity failures such as unequal device access can persist even where an institution's regulatory environment is otherwise favourable; folding equity into environment risks the downstream, secondary treatment both this paper's respondents and the equity literature warn against [13].

Table 6 positions the proposed framework against the three adoption models most used in the Kenyan and regional e learning literature reviewed. TAM-based studies centre on individual perceived usefulness and ease of use; UTAUT2-based studies add habit, price value and facilitating conditions but remain individual-level; TOE based studies operate at institutional level across technology, organisation and environment but were not developed for autonomous, decision-making systems and do not, without extension, treat data-governance accountability or learner equity as separate, named constructs [14], [15]. The proposed framework is deliberately institutional rather than individual (like TOE, unlike TAM/UTAUT2), and names two constructs, explicit AI accountability and cross-cutting equity, separately from TOE's environmental dimension, on the grounds that these become first-order, trackable concerns only once a system acts autonomously on learner data rather than merely being used by a learner.

Table 6. Prior Models Versus the Proposed Framework

Model	Level	Names AI accountability?	Names equity?
TAM	Individual	No	No
UTAUT2	Individual	No	No
TOE [14]	Institutional	Not separately	Not separately
Proposed (this study)	Institutional	Yes	Yes (cross-cutting)

Four-Domain Structure

The framework treats infrastructure, policy and governance, human capability, and equity as interdependent rather than sequential domains, summarised in Table 7: strength in one (e.g., extensive Moodle customisation) does not substitute for weakness in another (e.g., unclear data-governance accountability for autonomous decisions), and the equity domain is deliberately positioned as a cross-cutting check on the other three rather than a final add-on.

Table 7. Four-Domain Agentic-AI Readiness Framework

Domain	Core question	Illustrative indicator
Infrastructure	Can the technical base sustain autonomous, always-on agents?	Hosting/API capacity, connectivity, device access
Policy & Governance	Are decisions and data flows accountable?	Formal AI-in-teaching policy, DPA-aligned data governance, accountability framework
Human Capability	Can staff and students operate, oversee and trust agentic tools?	Staff AI-management skills, student digital readiness, leadership commitment
Equity	Does the transition avoid widening access gaps?	Rural/remote parity, disability accommodation, affordability

Operationally, the framework suggests a staged pathway: (1) baseline audit of the four domains using instruments similar to the one applied here, ideally at institutional rather than individual level; (2) closing governance gaps (a formal AI-in-teaching policy and an explicit accountability framework for AI-driven decisions) before scaling any pilot beyond integrity or routing tools; (3) targeted capacity-building for both staff and students, since this sample suggests neither group can be assumed to lag the other; and (4) an equity gate at each stage, requiring evidence that rural, remote and disability-related access gaps are being actively monitored and mitigated before a pilot is declared ready for full institutional deployment.

Practical Readiness-Audit Checklist

Table 8 translates the framework into a short checklist that a systems librarian or ICT/eLearning director can apply directly during an institutional audit, prior to approving any expansion of adaptive or agentic functionality beyond a contained pilot.

Table 8. Pre-Scaling Readiness Checklist

Element	Example value
Infrastructure	Hosting/API capacity load-tested; connectivity mapped by campus/region
Policy & Governance	AI-in-teaching policy adopted; named accountability owner per workflow
Human Capability	Staff AI-management training completed; student orientation delivered
Equity	Rural/remote and disability access gaps assessed and mitigation funded

Future Research Directions

Building on the shortcomings identified in the Limitation Section, four directions for future research follow. First, a larger, adequately powered replication should recruit a stratified, institution-identified sample balanced between public and private provision, so the individual- and institution-level analyses originally envisaged for this study can be run with appropriate power. Second, self-reported infrastructure and budget items should be paired with independently verified institutional data (audited ICT expenditure, measured connectivity) to test whether the benign infrastructural-barrier ratings observed here hold once triangulated and once the public-university share is increased. Third, since current adoption in this sample clusters around integrity and routing tools rather than autonomous agents, a follow-up study once genuinely agentic tools are deployed at scale would test whether these readiness patterns persist further along the agentic spectrum [1], [2], [4]. Fourth, the accountability-owner and scoped-permission governance mechanisms noted in the Discussion Section [3], [16] merit a dedicated case study within a Kenyan ODeL institution to assess their practicality under local resource constraints.

Conclusion and recommendation

This study set out to establish an empirical baseline for agentic AI readiness within Moodle-based ODeL delivery in Kenya. Within a modest, exploratory and private-institution-skewed sample of seventeen respondents, the picture that emerges is one of technically active but governance-cautious institutions: extensive Moodle customisation and a majority uptake of automated integrity and pathway tools, set against low visibility of ICT/AI budgets, self-identified gaps in formal AI policy and accountability, and persistent concern about rural and remote student disadvantage. Neither a staff-student readiness gap nor a budget-adoption association reached statistical significance in this sample, a finding that should motivate a larger, adequately powered, multi-

institution study rather than be read as a substantive conclusion.

For systems librarians, ICT directors and ODeL administrators, four recommendations follow directly from the data. First, formalise AI-in-teaching and data-governance policy ahead of, not after, scaling adaptive functionality, given that policy items scored no higher on average than infrastructure items despite Data Protection Act compliance being rated highly. Second, invest deliberately in equity safeguards for rural and remote learners before or alongside any agentic AI rollout, rather than treating equity as a downstream concern, since 75% of respondents in this sample agreed that such learners already face disadvantages in accessing AI-enhanced features. Third, commission institution-level, not only individual-level, readiness audits with verified rather than self-reported budget and infrastructure data, deliberately including public universities, to enable the kind of adequately powered inferential analysis this exploratory study could not support. Fourth, adapt the accountability-owner and scoped-permission governance practices summarised in the Discussion Section to the Kenyan ODeL context, so that responsibility for any agentic decision affecting a student remains traceable to a named individual even as functionality becomes more autonomous.

Systems librarians occupy a distinctive position from which to lead this transition: they already sit at the intersection of information governance, digital infrastructure and user support, and their existing responsibility for metadata, access control and information ethics maps closely onto the policy, human-capability and equity domains of the framework proposed in the proposed conceptual readiness framework. Positioning the library and ICT/eLearning functions jointly as institutional stewards of the readiness audit recommended above, rather than leaving agentic AI adoption to be driven by vendor offers or ad hoc departmental pilots, would give Kenyan ODeL institutions a coherent, accountable route from static LMS infrastructure toward agentic-AI-driven delivery.

Acknowledgment

The authors thank the institutional leaders, ICT/eLearning directors, ODeL coordinators and students who took part in the survey. The authors used Claude (Anthropic) to organise this manuscript into the required IEEE conference structure.

Declarations

Conflicts of interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] M. A. Raidas and R. Bhandari, "Agentic AI in Education: Redefining Learning for the Digital Era," in Artificial Intelligence in Education: Transforming Learning for the Future, London, U.K.: Scholars' Press, 2025, p. 89.

[2] G. Kostopoulos, V. Gkamas, M. Rigou, and S. Kotsiantis, "Agentic AI in Education: State of the Art and Future Directions," IEEE Access, vol. 13, pp. 177467–177491, 2025.

[3] I. Adabara, B. O. Sadiq, A. N. Shuaibu, Y. I. Danjuma, and V. Maninti, "Trustworthy Agentic AI Systems: A Cross-Layer Review of Architectures, Threat Models, and Governance Strategies for Real-World Deployment," F1000Research, vol. 14, p. 905, 2025.

[4] L. Yan, "From Passive Tool to Socio-cognitive Teammate: A Conceptual Framework for Agentic AI in Human-AI Collaborative Learning," arXiv:2508.14825, 2025.

[5] S. Murugesan, "The Rise of Agentic AI: Implications, Concerns, and the Path Forward," IEEE Intelligent Systems, vol. 40, no. 2, pp. 8–14, Mar./Apr. 2025.

[6] N. Kshetri, "Revolutionizing Higher Education: The Impact of Artificial Intelligence Agents and Agentic Artificial Intelligence on Teaching and Operations," IT Professional, vol. 27, 2025.

[7] R. N. Kibuku, D. O. Ochieng, and A. N. Wausi, "E-Learning Challenges Faced by Universities in Kenya: A Literature Review," Electronic Journal of e-Learning, vol. 18, no. 2, pp. 150–161, 2020.

[8] A. S. Mulwa and D. N. Kyalo, "The Influence of Institutional Factors on the Readiness to Adopt E-Learning in Kenyan Universities," Journal of Education and Practice, vol. 2, no. 11, pp. 1–9, 2011.

[9] M. Kashorda and T. Waema, E-Readiness Survey of Kenyan Universities Report, Nairobi, Kenya: Kenya Education Network (KENET), 2014.

[10] J. W. Mwangi, "Challenges Facing the Implementation of E-Learning in Kenyan Universities," International Journal of Academic Research in Progressive Education and Development, vol. 5, no. 3, pp. 1–12, 2016.

[11] Republic of Kenya, The Data Protection Act, No. 24 of 2019, Kenya Gazette, vol. CXXI, no. 156, 15 Nov. 2019.

[12] Office of the Data Protection Commissioner, Guidance Note for the Education Sector, Nairobi, Kenya: ODPC Kenya, 2024.

[13] R. Lembani, A. Gunter, M. Breines, and M. T. B. Dalu, "The Same Course, Different Access: The Digital Divide Between Urban and Rural Distance Education Students in South Africa," Journal of Geography in Higher Education, vol. 44, no. 1, pp. 70–84, 2020.

[14] L. G. Tornatzky and M. Fleischer, The Processes of Technological Innovation. Lexington, MA: Lexington Books, 1990.

[15] J.-E. Kalmus and A. Nikiforova, "To Accept or Not to Accept? An IRT TOE Framework to Understand Educators' Resistance to Generative AI in Higher Education," in Proc. Int. Workshop on AI in Education and Educational Research (AIEER), co-located with ECAI 2024, arXiv:2407.20130, 2024.

[16] K. Navaie, "From Rights to Runtime: Privacy Engineering for Agentic AI," AI Magazine, vol. 46, no. 4, e70036, 2025.

[17] A. Azevedo et al., "Institutional Policies on Artificial Intelligence in Higher Education: Frameworks and Best Practices for Faculty," New Directions for Adult and Continuing Education, Hoboken, NJ: Wiley, 2025.