

RESEARCH TITLE

Artificial Intelligence Literacy and Responsible Use in Sudanese Higher Education: A Narrative Review and Practical Framework

Balla Mustafa Idris Elias¹

¹ Faculty of Languages, English Language Department, International University of Africa, Khartoum, Sudan
Correspondence: balla.idris65@gmail.com

HNSJ, 2026, 7(9); <https://doi.org/10.53796/hnsj79/56>

Received at 15/08/2026

Accepted at 22/08/2026

Published at 01/09/2026

Abstract

Artificial intelligence (AI), particularly generative AI, is entering higher education faster than institutions can develop shared rules for its use. This narrative review examines what AI systems can and cannot do in teaching, learning, assessment, and academic work, and translates the evidence into a practical framework for Sudanese universities. A purposive search of scholarly databases and authoritative organizational repositories was conducted in September 2026. Sources were selected for their relevance to AI concepts, higher education, generative-AI risks, AI literacy, academic integrity, Arabic-language technologies, and responsible governance. The literature indicates that AI can support feedback, content adaptation, language assistance, and administrative efficiency, but demonstrated benefits depend on task design, data quality, teacher expertise, and evaluation. Generative systems may produce fluent but false content, reproduce bias, expose personal data, and complicate assessment. Evidence specific to Sudanese higher education and Sudanese Arabic remains limited; therefore, local effectiveness should not be inferred from studies conducted elsewhere. The article proposes a staged framework centred on governance, human oversight, assessment redesign, staff and student AI literacy, Arabic and local-language evaluation, equitable access, and continuous monitoring. AI should be treated as a fallible support tool rather than an autonomous authority. The principal contribution is a context-sensitive agenda that separates established evidence from plausible applications and identifies priorities for local empirical research.

Key Words: artificial intelligence; generative AI; AI literacy; higher education; academic integrity; Sudan; Arabic language.

المعرفة بالذكاء الاصطناعي والاستخدام المسؤول له في التعليم العالي السوداني: مراجعة سردية وإطار عملي

المستخلص

يشهد التعليم العالي انتشاراً سريعاً لتطبيقات الذكاء الاصطناعي، ولا سيما الذكاء الاصطناعي التوليدي، في وقت لا تزال فيه كثير من المؤسسات تعمل على وضع سياسات واضحة لاستخدامه. تهدف هذه المراجعة السردية إلى توضيح قدرات أنظمة الذكاء الاصطناعي وحدودها في التدريس والتعلم والتقييم والعمل الأكاديمي، وتحويل الأدلة المتاحة إلى إطار عملي ملائم للجامعات السودانية. أُجري في سبتمبر 2026 بحث هادف في قواعد بيانات علمية ومستودعات منظمات دولية، واختيرت المصادر ذات الصلة بالمفاهيم الأساسية، والتعليم العالي، ومخاطر الذكاء الاصطناعي التوليدي، والثقافة بالذكاء الاصطناعي، والنزاهة الأكاديمية، وتقنيات اللغة العربية، والحوكمة المسؤولة. تشير الأدبيات إلى إمكان دعم التغذية الراجعة، وتكييف المحتوى، والمساعدة اللغوية، والكفاءة الإدارية، إلا أن الفوائد المثبتة تعتمد على تصميم المهمة وجودة البيانات وخبرة المعلم وأساليب التقييم. وقد تنتج الأنظمة التوليدية محتوى مقنعاً لكنه غير صحيح، كما قد تعيد إنتاج التحيز، وتعرض البيانات الشخصية للخطر، وتعدّد التحقق من أصالة أعمال الطلاب. ولا تزال الأدلة الخاصة بالتعليم العالي السوداني واللهجة السودانية محدودة؛ لذلك لا يجوز تعميم نتائج البيئات الأخرى عليها دون اختبار محلي. وتقتصر المقالة إطاراً مرحلياً يركز على الحوكمة، والإشراف البشري، وإعادة تصميم التقييم، وتنمية ثقافة الذكاء الاصطناعي لدى الطلاب والعاملين، وتقييم الأداء باللغة العربية واللهجات المحلية، والعدالة في الوصول، والمتابعة المستمرة. وتتمثل مساهمة الدراسة في تقديم أجندة تراعي السياق، وتفصل بين الأدلة القائمة والتطبيقات المحتملة، وتحدد أولويات للبحث التجريبي المحلي.

الكلمات المفتاحية: الذكاء الاصطناعي؛ الذكاء الاصطناعي التوليدي؛ الثقافة بالذكاء الاصطناعي؛ التعليم العالي؛ النزاهة الأكاديمية؛ السودان؛ اللغة العربية.

1. Introduction

Artificial intelligence has become a general-purpose set of computational methods used in education, health, finance, communication, and public administration (Dwivedi et al., 2021). In higher education, recent attention has centred on generative AI systems that produce text, images, audio, or code in response to prompts. Their accessible interfaces create opportunities for tutoring, feedback, translation, and content development, while also raising questions about reliability, privacy, fairness, authorship, and the purposes of assessment (Kasneci et al., 2023; Miao & Holmes, 2023). These questions are especially important where institutions face uneven connectivity, limited budgets, multilingual classrooms, or incomplete policy frameworks.

The term AI must not be equated with a single company, product, or chatbot. AI is the wider field concerned with computational systems that perform functions commonly associated with perception, language, learning, prediction, planning, or decision support (Russell & Norvig, 2021). Machine learning is one approach within AI; deep learning is a family of machine-learning methods based on multilayer neural networks; and generative AI refers to models designed to produce new content. These distinctions matter because the capabilities and risks of a medical image classifier, a recommender system, and a large language model are not identical.

This article is a narrative review, not an empirical study or a systematic review. Its objective is to synthesize selected conceptual, educational, and ethical literature and to propose a practical framework for responsible AI use in Sudanese higher education. Its intended audience includes university leaders, lecturers, curriculum designers, quality-assurance units, librarians, and students. The review focuses on teaching, learning, assessment, and academic work; it does not evaluate a specific product, measure adoption in Sudan, or estimate the labour-market effects of AI. The contribution is threefold: it corrects common technical misconceptions, separates possible uses from demonstrated outcomes, and identifies actions and research priorities appropriate to a context in which local evidence remains limited.

2. Review Approach

A purposive narrative search was undertaken in September 2026 using scholarly search services, publisher databases, the ACM Digital Library, ACL Anthology, and UNESCO's document repository. Combinations of the following terms were used: artificial intelligence, machine learning, deep learning, generative AI, large language models, higher education, teaching, assessment, academic integrity, AI literacy, hallucination, bias, privacy, Arabic, dialectal Arabic, Sudan, and responsible AI. Reference lists of central publications were also examined to identify foundational works.

Priority was given to peer-reviewed reviews, conceptual papers, benchmark or technical studies, and authoritative international guidance published from 2016 to 2026. Earlier or book-length sources were retained when they provided foundational definitions. Sources were included when they directly informed at least one of the review questions: What capabilities are relevant to higher education? What educational benefits and risks have been reported? What competencies and governance safeguards are recommended? What evidence is available for Arabic varieties and the Sudanese context? Product marketing, duplicated commentary, and claims without an identifiable source were excluded. The evidence was compared by setting, task, type of publication, outcome addressed, and stated limitation, then organized thematically.

This approach is transparent but not exhaustive. It did not use a registered protocol, duplicate screening, formal risk-of-bias scoring, or meta-analysis. English-language material predominated, and rapidly changing models can make product-specific findings obsolete. The

scarcity of directly relevant Sudanese studies is itself an important limitation: the framework offered below is an evidence-informed proposal that requires local testing, not proof of effectiveness in Sudan.

3. Conceptual Foundations

AI systems transform inputs into outputs according to computational rules and learned parameters. Some systems use hand-crafted rules; many contemporary systems learn statistical patterns from data. Their outputs may be classifications, rankings, forecasts, recommendations, generated content, or actions. Performance is therefore task-specific and must be evaluated against an appropriate benchmark. No model predicts new cases perfectly, and high average accuracy can conceal serious errors for particular groups or rare cases.

Concept	Working definition	Educational example
Artificial intelligence	The broad field of computational systems performing tasks associated with perception, language, learning, prediction, planning, or decision support.	An adaptive tutoring platform or automated scheduling system.
Machine learning	Methods that estimate patterns or decision functions from data rather than relying only on explicitly programmed rules.	A model that predicts which students may need additional support.
Deep learning	Machine learning using multilayer neural networks that learn representations from large or complex datasets.	Speech recognition or image classification.
Generative AI	Models that generate new text, images, audio, video, or code based on learned statistical patterns and user input.	Drafting practice questions or explaining a concept in different ways.
Large language model	A generative model trained on large text collections to estimate and generate sequences of language units. Next-token prediction is central to many LLMs, but it does not describe all AI systems.	A conversational assistant used to brainstorm or revise prose.

Table 1. Core concepts used in this review. Definitions are synthesized from Russell and Norvig (2021) and Goodfellow, Bengio, and Courville (2016).

Neural networks are mathematical models inspired in part by ideas from neuroscience, but they are not literal copies of the human brain (Goodfellow et al., 2016). Likewise, fluent language output should not be treated as evidence of consciousness or guaranteed reasoning. For large language models, factual accuracy, reasoning performance, interpretability, and consciousness are separate questions. A model may produce a correct answer through unreliable steps, an incorrect answer in persuasive language, or an output whose internal basis is difficult to interpret.

Catastrophic forgetting is another frequently misunderstood concept. It refers to the loss of performance on earlier tasks when a neural network is trained sequentially on new tasks. It is a model-training problem, not a claim that every conversation with a user erases earlier knowledge. Kirkpatrick et al. (2017) demonstrated elastic weight consolidation as one mitigation strategy for sequential learning.

4. Uses of AI in Higher Education

4.1 Teaching and learning support

AI can assist lecturers with examples, practice questions, rubrics, lesson outlines, and alternative explanations. It can also help students brainstorm, receive formative feedback, practice a language, or obtain an initial explanation outside class time. Earlier work on AI in higher education included profiling, prediction, adaptive systems, intelligent tutoring, and assessment; however, a systematic review found that educators' perspectives were often underrepresented (Zawacki-Richter et al., 2019). Generative AI has expanded access to conversational support, but accessibility does not establish pedagogical effectiveness. Learning gains depend on alignment with objectives, the quality of interaction, prior knowledge, and teacher guidance (Kasneci et al., 2023; Tlili et al., 2023).

The most defensible use is often augmentation: AI produces a draft, suggestion, or explanation that a learner or expert must examine. For example, a lecturer may generate several quiz items and then verify their factual accuracy, cognitive level, cultural suitability, and language. A student may compare an AI explanation with course readings and document the corrections made. Such designs keep the intellectual work visible and make verification part of learning.

4.2 Assessment and feedback

AI may accelerate low-stakes feedback and help create varied practice activities, but it can also generate work that obscures what a student knows. The educational problem is not adequately described by labelling all users as cheaters. Misconduct depends on the rules of the assessment, the student's disclosure, and whether the submitted work represents the required learning. Institutions should define permitted, restricted, and prohibited uses for each task. Assessment can be strengthened through oral defence, supervised components, staged drafts, process logs, local or personal application, and reflective explanation of any AI assistance. Reliance on automated AI-text detectors should be cautious because detection is probabilistic and can create false accusations.

4.3 Research, writing, and administration

Generative tools may support search-term development, language editing, coding assistance, summarization of text supplied by the user, and preparation of routine communications. They must not replace reading of primary sources or responsibility for claims. Generated references may be nonexistent or inaccurate, and confidential research data should not be entered into an external service without authorization. In administration, predictive or ranking systems may support scheduling and service delivery, but decisions

affecting admission, progression, employment, or discipline require documented criteria, data protection, bias testing, and a meaningful route for human review.

Use	Plausible value	Main risk	Minimum safeguard
Drafting learning materials	Faster generation of examples and variants	Errors, bias, or poor curricular alignment	Lecturer review against learning outcomes and sources
Student tutoring	Immediate explanations and practice	Over-reliance or misleading feedback	Verification tasks, reflection, and access to a teacher
Language support	Revision, translation, and practice	Loss of voice or inaccurate terminology	Bilingual review and disclosure of substantial assistance
Assessment	Question generation and formative feedback	Authorship ambiguity and invalid inference about learning	Clear rules, process evidence, oral or supervised components
Research support	Brainstorming and coding assistance	Fabricated citations, confidentiality, hidden errors	Primary-source checking, data controls, author accountability
Administrative decisions	Pattern detection and workflow support	Discrimination, opacity, and automation bias	Impact assessment, audit, explanation, and human appeal

Table 2. Potential higher-education uses, risks, and safeguards. Values are conditional rather than guaranteed outcomes.

5. Risks, Limitations, and Ethical Concerns

5.1 Factual reliability and reasoning

Generative systems can produce statements that are fluent, specific, and false. The literature commonly calls this phenomenon hallucination, although the term should not imply human perception. A broad survey shows that factual inconsistency arises from data, modelling, inference, and evaluation choices and remains difficult to eliminate completely (Ji et al., 2023). Risk is higher when users cannot independently verify the answer or when the task concerns medicine, law, finance, safety, or academic evidence. Confidence of expression is not a measure of truth.

5.2 Bias, language, and cultural fit

Models can reproduce social stereotypes and uneven performance present in their training data or evaluation design. Large training corpora may also overrepresent dominant languages and viewpoints (Bender et al., 2021). Arabic is not a single uniform test condition: Modern Standard Arabic differs from national and regional varieties, and performance varies by task, dataset, domain, and model version. Arabic-specific models such as ARBERT and MARBERT demonstrate the value of distinguishing formal and dialectal data rather than treating all Arabic as one category (Abdul-Mageed et al., 2021).

This review located no sufficiently robust, peer-reviewed benchmark that would justify a general numerical claim about generative-AI performance in Sudanese Arabic across educational tasks. The responsible conclusion is therefore not that AI is necessarily weak in all Sudanese or Standard Arabic uses. It is that local performance is uncertain and should be measured. Universities should test representative prompts, subjects, terminology, and student language practices before adopting a tool at scale, and they should record the model and version because performance can change.

5.3 Privacy, transparency, and accountability

Prompts may contain student records, unpublished research, examination material, or sensitive personal information. Institutions need rules on what data may be entered, where data are stored, whether prompts are used for training, and how long records are retained. Transparency also requires users to know when AI materially shaped content or a decision. Yet disclosure alone is insufficient: responsibility must remain with identifiable people who can examine evidence, correct errors, and respond to appeals. UNESCO's human-centred guidance emphasizes protection of human agency, inclusion, privacy, and institutional capacity (Miao & Holmes, 2023; UNESCO, 2021).

5.4 Academic integrity and cognitive engagement

Generative AI can be used to bypass learning, but categorical statements such as 'students become cheaters' or 'they will never learn' are unsupported. The effect depends on task design and use. A tool that supplies a final answer with no evaluation may reduce productive effort; the same tool can prompt comparison, critique, revision, or explanation. Academic-integrity research therefore recommends explicit expectations, assessment redesign, and education rather than relying only on prohibition (Cotton et al., 2024). AI literacy should include the ability to recognize AI, understand basic capabilities and limitations, evaluate outputs, communicate effectively with systems, and consider ethical consequences (Long & Magerko, 2020).

5.5 Access, cost, and sustainability

Effective use may require reliable electricity, internet access, suitable devices, paid subscriptions, and staff time. These requirements can widen existing inequalities when some learners receive continuous access and others do not. Large-scale model training and operation also consume computational and energy resources. A university should therefore evaluate educational value against total cost, accessibility, data governance, and environmental implications rather than adopting a system because it is novel. Low-bandwidth alternatives, campus access, accessible interfaces, and non-AI routes to complete assessed work are important equity measures.

6. Implications for Sudanese Higher Education

Sudanese universities vary in discipline, resources, language of instruction, and student circumstances. A single national or institutional rule is unlikely to address every assessment. Nevertheless, a common governance baseline can protect students and staff while allowing responsible experimentation. Because direct local evidence is limited, implementation should begin with small, documented pilots and predefined educational outcomes. Claims of improved learning, reduced workload, or lower cost should be tested rather than assumed.

Language policy deserves particular attention. Institutions may use English, Modern Standard Arabic, Sudanese Arabic, or combinations of these in teaching and student support. Evaluation should therefore distinguish language varieties and tasks. Accuracy in conversational explanation does not establish accuracy in academic translation, discipline-

specific terminology, automated scoring, or speech recognition. Local academics, language specialists, students, and disability-support staff should participate in evaluation.

7. A Practical Framework for Responsible Adoption

Domain	Immediate action	Evidence to collect	Decision criterion
Governance	Create an accountable cross-functional committee and publish approved, restricted, and prohibited uses.	Incident records, user feedback, data flows, vendor terms	Risks have owners, controls, and appeal routes.
Pedagogy	Approve small pilots tied to explicit learning outcomes.	Learning evidence, teacher workload, student experience	Benefit exceeds non-AI alternatives for the defined task.
Assessment	State AI rules for each task and require evidence of process where relevant.	Validity, misconduct patterns, false accusations, accessibility	Assessment still measures the intended learning.
AI literacy	Train staff and students in prompting, verification, citation, privacy, bias, and limits.	Pre/post competency tasks and error-detection performance	Users can verify and disclose use responsibly.
Language	Benchmark named models on MSA, Sudanese Arabic, English, and disciplinary terminology as needed.	Task-level accuracy and qualitative error analysis	Performance is acceptable for the exact population and purpose.
Equity and access	Provide low-bandwidth and non-AI alternatives, accessible formats, and transparent cost planning.	Access rates, device/connectivity barriers, disability impacts	No group is unfairly excluded from learning or assessment.
Monitoring	Review tools, policies, and model versions each academic term.	Errors, complaints, learning outcomes, cost, policy compliance	Continue, modify, or stop on documented evidence.

Table 3. Proposed staged framework for Sudanese higher-education institutions.

The framework begins with governance because educational experimentation without clear responsibility can expose learners to avoidable harm. Policies should name permitted purposes, prohibited data, disclosure expectations, procurement checks, and procedures for

complaints. They should be understandable to students and revised as technologies and evidence change. High-impact decisions should never depend solely on an opaque automated output.

The pedagogical principle is constructive use: students should do intellectual work that remains observable. Appropriate activities include checking an AI answer against assigned sources, identifying unsupported assumptions, comparing outputs in Arabic and English, revising a weak draft with an explanation of changes, and testing the same prompt across systems. These practices make AI literacy part of disciplinary learning rather than a separate technical topic.

8. Research Agenda

A publication agenda for Sudan should move beyond general attitudes toward task-specific evidence. Priority studies include:

- baseline surveys and performance-based measures of staff and student AI literacy across disciplines and institutions;
- controlled or quasi-experimental studies comparing AI-supported and non-AI learning activities using valid educational outcomes;
- benchmarks for Modern Standard Arabic, Sudanese Arabic, English-Arabic academic translation, and discipline-specific terminology, with named model versions;
- qualitative studies of student and lecturer practices, including disclosure, verification, trust, and barriers to access;
- audits of privacy, bias, accessibility, and total cost in institutional pilots; and
- assessment studies examining oral defence, staged drafting, process portfolios, and authentic local tasks.

Future studies should report the model, version, date of access, prompt design, language variety, task, comparison condition, and error criteria. Without these details, results are difficult to reproduce and may become misleading after a system update. Local evidence should also avoid treating Sudanese students or institutions as a homogeneous group.

9. Conclusion

AI is a broad field, and generative AI is only one part of it. In higher education, these technologies can support teaching, learning, writing, and administration, but their outputs are not inherently accurate, fair, private, or educationally effective. Benefits reported in one task or setting cannot be generalized automatically to another. Responsible adoption therefore requires human oversight, explicit assessment rules, verification of sources, protection of personal data, equitable access, and evaluation of named systems for defined purposes.

For Sudanese higher education, the immediate need is not an unrestricted embrace or a categorical ban. It is a staged programme of AI literacy, governance, local language evaluation, and empirical research. The absence of robust evidence for Sudanese Arabic and local university contexts should be presented as a research gap, not converted into an unsupported claim. AI can be a useful assistant when people retain responsibility for objectives, judgment, and consequences.

Declarations

Funding: The author received no specific funding for this work.

Conflict of interest: The author declares no conflict of interest.

Ethics statement: This narrative review did not involve human participants or identifiable personal data.

Data availability: No new dataset was generated or analysed for this article.

References

- Abdul-Mageed, M., Elmadany, A., & Nagoudi, E. M. B. (2021). ARBERT & MARBERT: Deep bidirectional transformers for Arabic. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (pp. 7088-7105). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.acl-long.551>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610-623). Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>
- Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 61(2), 228-239. <https://doi.org/10.1080/14703297.2023.2190148>
- Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P. V., Janssen, M., Jones, P., Kar, A. K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., ... Williams, M. D. (2021). Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57, Article 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press. <https://www.deeplearningbook.org/>
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), Article 248. <https://doi.org/10.1145/3571730>
- Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, Article 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521-3526. <https://doi.org/10.1073/pnas.1611835114>
- Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1-16). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376727>
- Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000386693>
- Russell, S. J., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
- Tlili, A., Shehata, B., Adarkwah, M. A., Bozkurt, A., Hickey, D. T., Huang, R., & Agyemang,

B. (2023). What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education. *Smart Learning Environments*, 10, Article 15. <https://doi.org/10.1186/s40561-023-00237-x>

UNESCO. (2021). Recommendation on the ethics of artificial intelligence. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>

Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education - Where are the educators? *International Journal of Educational Technology in Higher Education*, 16, Article 39. <https://doi.org/10.1186/s41239-019-0171-0>