RITAM ISH PUBLICATIONS

Journal of Progress in Artificial Intelligence
https://www.jpairitamish.org/

ISSN: 3139-4957 (Online)

Vol. 2, Iss. 3, July 2026 – September 2026

From AI assistance to responsible academic agency: A human-centered framework for generative artificial intelligence in higher education

Benjamin Batte[1]

[1]Ph.D. Student in Information Technology (Artificial Intelligence Concentration), University of the Cumberlands, Williamsburg, Kentucky, USA

Received: 03-Aug-2026, Manuscript No. JPAI-2026-0022; **Editor assigned:** 07-Aug-2026, PreQC No. JPAI-2026-0022 (PQ); **Reviewed:** 24-Aug-2026, QC No. JPAI-2026-0022; **Revised:** 27-Aug-2026, Manuscript No. JPAI-2026-0022 (R); **Published:** 31-Aug-2026

Citation: Batte, B. (2026). From AI assistance to responsible academic agency: A human-centered framework for generative artificial intelligence in higher education. J Prog Artif Intell.

Corresponding: Benjamin Batte, Email: benjamin.kamulali@gmail.com

ABSTRACT

Generative artificial intelligence has moved from a peripheral educational technology to an ordinary part of university study, writing, research, and assessment. Existing scholarship has documented substantial opportunities in personalised support, accessibility, feedback, and productivity, while also identifying persistent concerns involving misinformation, academic misconduct, privacy, bias, inequitable access, and cognitive dependence. However, the debate is too often organised around a narrow question: whether artificial intelligence should be allowed or prohibited. This conceptual article argues that the more consequential question is whether students retain responsible academic agency when artificial intelligence participates in their learning. Drawing on research in artificial intelligence in education, self-regulated learning, academic integrity, human-centred artificial intelligence, and institutional governance, the article develops a Responsible Academic Agency Framework. The framework integrates six interdependent domains: purpose and learning alignment; cognitive ownership and human judgment; epistemic verification; transparency and attribution; equity and privacy; and institutional governance and assessment validity. The central proposition is that artificial intelligence contributes to education only when it strengthens rather than substitutes the learner's capacity to set goals, evaluate evidence, make decisions, explain conclusions, and accept responsibility for academic work. The framework provides practical guidance for students, educators, university leaders, and technology providers, and offers a research agenda for evaluating learning quality rather than focusing solely on adoption. The article contributes to the literature by shifting the discussion from tool access and rule compliance toward the preservation, demonstration, and development of human agency within artificial intelligence-mediated higher education.

KEYWORDS: Generative artificial intelligence, higher education, academic agency, academic integrity, AI, literacy, assessment, governance, human-centred AI

1.INTRODUCTION

Generative artificial intelligence has altered the conditions under which university students read, write, research, solve problems, and demonstrate learning. Tools based on large language models can produce explanations, outlines, summaries, code, translations, feedback, and extended academic prose within seconds. Their accessibility has accelerated student adoption and challenged universities to reconsider long-established assumptions about authorship, independent work, assessment validity, and the meaning of academic competence. Artificial intelligence is, therefore, not simply another classroom application. It is a sociotechnical change that redistributes cognitive work among students, educators, institutions, and computational systems. Research on artificial intelligence in higher education has expanded rapidly. A meta-systematic review documented applications in student profiling, prediction, tutoring, assessment, feedback, and administrative decision making [1]. Earlier systematic research found that higher education applications were concentrated in profiling and prediction, intelligent tutoring systems, assessment, and adaptive support [2]. A broad review of the field confirmed rapid growth in teaching, learning, and student support applications while identifying persistent gaps in pedagogical and ethical analysis [3]. Related analyses reach a similar conclusion, showing that artificial intelligence carries simultaneous ethical, social, and educational implications for universities [4], and empirical work confirms measurable effects of artificial intelligence adoption on teaching and learning outcomes [5].Broader commentary on large language models in education has offered a similarly wide-angle account of these developments, cataloguing opportunities such as personalised tutoring and feedback alongside risks, including overreliance, misinformation, and inequitable access. Recent scholarship has focused specifically on generative artificial intelligence and the risks it creates for academic integrity, including unauthorized assistance, uncertain authorship, and overreliance on detection tools [6]. Other research emphasizes the need to balance innovation with academic standards as universities integrate generative artificial intelligence into teaching and assessment [7]. Ethical scholarship has highlighted fairness, privacy, agency, transparency, and human oversight as essential conditions for responsible use in education [8]. Complementary scholarship on large language models similarly outlines both the pedagogical opportunities and the ethical challenges these systems introduce for education [9]. A systematic review of responsible artificial intelligence in education further identifies bias, privacy loss, opacity, and unequal access as risks that require mitigation throughout the technology lifecycle [10]. This concern about social and educational consequences echoes a broader ethical stock-take of artificial intelligence adoption across university functions, which likewise calls for deliberate attention to fairness and institutional readiness rather than incidental adoption. Universities are also developing policy frameworks that connect pedagogy, governance, and institutional operations [11]. Global policy analysis shows growing adoption of institutional guidance, training programs, and authentic assessment practices, although gaps remain in data governance, equity, and consistent communication [12]. Despite this expanding literature, the dominant debate often remains trapped between adoption and prohibition. In one position, artificial intelligence is presented as an inevitable productivity technology that universities must embrace to prepare students for employment. In the opposing position, it is treated primarily as a threat to learning and academic integrity. A more optimistic strand of this literature argues that, when thoughtfully integrated, artificial intelligence can expand access, enrich pedagogy, and strengthen rather than undermine the broader mission of higher education. Both positions can obscure the central educational issue. A university does not fulfill its purpose merely because students' complete tasks efficiently, nor does it protect learning merely by restricting a tool. The central question is whether the learning process continues to develop human judgment, disciplinary understanding, intellectual responsibility, and the capacity to explain and defend one's work. Some scholars have advanced an optimistic reading of this shift, arguing that artificial intelligence can expand access and opportunity across higher education when adoption is managed thoughtfully [13].

This article introduces responsible academic agency as a unifying concept for evaluating generative artificial intelligence in higher education. Academic agency refers to the learner's capacity to form intentions, make informed choices, regulate effort, evaluate evidence, revise understanding, and accept responsibility for academic decisions. Social cognitive theory emphasizes intentionality, forethought, self-regulation, and self-reflection as central features of human agency [14]. Self-regulated learning similarly explains effective learning as an active process in which students set goals, select strategies, monitor progress, and evaluate outcomes [15]. A responsible academic agency adds ethical and epistemic obligations. Students must not only exercise choice, but must also verify claims, disclose material assistance, protect sensitive information, respect academic rules, and remain accountable for the final work. The article makes three contributions. First, it reframes the generative artificial intelligence debate from a question of tool permission to a question of human capability and responsibility. Second, it proposes a six-domain Responsible Academic Agency Framework that links student practice, pedagogy, assessment, ethics, and institutional governance. Third, it develops a research agenda that encourages empirical studies to examine learning quality, judgment, verification behavior, and equity rather than relying mainly on adoption rates and attitudes. The resulting framework is intended to be sufficiently general for institutional policy while remaining concrete enough to guide classroom practice.

2.CONCEPTUAL FOUNDATIONS

2.1 Generative Artificial Intelligence as a Sociotechnical Educational Change

Artificial intelligence in education is frequently discussed as if educational outcomes were properties of the technology itself. This view is incomplete. The same tool may support learning in one context and weaken it in another because outcomes depend on the interaction among task design, learner knowledge, institutional rules, access conditions, teacher guidance, and system capabilities. Reviews show that benefits are strongest when technologies are embedded within purposeful pedagogies rather than used as isolated replacements for instruction [1]. Research on student engagement similarly indicates that artificial intelligence is most effective when combined with interactive teaching, scaffolded feedback, project work, and active learning [16]. A chatbot that generates an answer can reduce productive struggle, while the same system can deepen learning when it asks diagnostic questions, offers counterexamples, or supports iterative feedback. A sociotechnical perspective also clarifies why policy statements alone are insufficient. Rules operate within cultures, assessment systems, technical infrastructures, and relationships of trust. An ecological policy framework identifies pedagogical, governance, and operational dimensions that universities must address together [11]. Global policy analysis also demonstrates that successful adoption requires clear responsibilities, communication, training, and continuing evaluation [12]. These findings suggest that responsible use cannot be delegated to individual students. It must be designed across the educational system.

2.2 Human Agency and Self-Regulated Learning

Human agency provides a strong theoretical foundation for evaluating artificial intelligence-mediated learning. Agency includes intentional action, anticipation of consequences, self-regulation, and self-reflection [14]. Effective learning is also a cyclical process in which learners set goals, select strategies, monitor progress, and evaluate outcomes [15]. Generative artificial intelligence can participate in every stage of this cycle. It can help define a problem, suggest strategies, monitor performance, and provide feedback. However, the educational value of this participation depends on whether the learner remains the author of goals and judgments. The danger is not simply that a machine produces text. The greater danger is the gradual transfer of intellectual control. When students accept generated answers without understanding, they may complete visible tasks while losing opportunities to develop invisible capabilities. These capabilities include uncertainty management, source evaluation, conceptual integration, disciplinary reasoning, and persistence. Empirical work shows that student use is heterogeneous. Some students use artificial intelligence to clarify concepts and improve drafts, while others use it primarily to minimize effort [17]. A responsible academic agency, therefore, cannot be inferred from the presence or absence of a tool. It must be evaluated through the student's decisions, explanations, verification practices, and demonstrated understanding.

2.3 Academic Integrity as a Learning Process

Academic integrity is often reduced to detecting prohibited outputs. Generative artificial intelligence makes this approach increasingly fragile, as generated text may not match existing sources, and detection systems can produce uncertain or unfair classifications. A stronger approach treats integrity as a process involving honesty, transparency, authorship, evidence, and accountability. Systematic reviews indicate that generative artificial intelligence may increase the risk of misconduct. Still, they also indicate that literacy, assessment redesign, and clear guidance are more sustainable responses than reliance on detection alone [6]. This conclusion echoes earlier commentary published soon after ChatGPT's release, which likewise argued that safeguarding integrity would depend on renewed pedagogical design rather than on detection technology alone [19]. Case-based analysis of chatbot use in education similarly cautions that the same conversational system can either scaffold or short-circuit learning depending on how a task is designed and how a student is guided to use it [18]. Process-based integrity requires students to understand which forms of assistance are permitted, document significant contributions made by artificial intelligence, verify information, and remain responsible for the submitted work. The Artificial Intelligence Disclosure Framework demonstrates how structured disclosure can make assistance visible without assuming that every use is equivalent [18]. Transparency is important, but disclosure by itself does not prove learning. A student may accurately disclose extensive artificial intelligence generation and still fail to demonstrate disciplinary competence. Integrity, therefore, needs to be joined with assessment validity and academic agency.

3. Educational Opportunities and their Conditions

Personalized academic support is one of the most frequently cited benefits of generative artificial intelligence. Students can request alternative explanations, examples, practice questions, translations, and feedback at times when a lecturer or peer is unavailable. This support may be particularly valuable for adult learners, multilingual students, students in large classes, and those balancing study with employment and family responsibilities. Artificial intelligence can reduce the delay between confusion and assistance, allowing students to continue learning rather than abandoning a task. Practical guidance developed for classroom use illustrates how instructors can adapt tools such as ChatGPT to concrete pedagogical practices, including scaffolded prompting, worked examples, and staged feedback, rather than leaving students to encounter the technology unsupported [19]. However, personalization should not be confused with automatic improvement. An explanation is educationally useful only when it is accurate, appropriately challenging, and connected to the intended learning outcome. Excessive simplification may produce confidence without competence. Similarly, immediate answers may interrupt the productive effort through which students learn to define problems and test possible solutions. The relevant design question is not whether assistance is personalized, but whether it supports the learner's next act of reasoning. Generative artificial intelligence can also support formative assessment. Students may compare an initial answer with a generated critique, request counterarguments, or use a model to identify unclear reasoning. Educators can create examples at different levels of quality and ask students to diagnose weaknesses. These uses shift artificial intelligence from answer production toward metacognitive support. Research on student engagement indicates that artificial intelligence contributes most effectively when combined with interactive teaching, scaffolded feedback, project work, and active learning [16]. Qualitative research with students and faculty reports a comparable mix of benefits and concerns, with participants describing gains in efficiency and personalization alongside worries about overreliance and diminished critical thinking [21]. Survey evidence from university settings points in the same direction, with respondents reporting both productivity gains and unresolved concerns about academic rigor and fairness [22]. Accessibility is another important opportunity. Text transformation, speech support, captioning, translation, and flexible explanation formats may reduce barriers for students with disabilities or language challenges. However, accessibility benefits depend on reliable access, inclusive design, and human support. If advanced tools are available only through costly subscriptions, the same technology can widen educational inequality. Institutions should therefore treat access as a governance issue rather than as an individual purchasing decision. Finally, universities have a responsibility to prepare graduates for workplaces in which artificial intelligence is increasingly used for communication, analysis, software development, research, and decision support. Preparation should extend beyond prompt production. Students need to learn how to formulate problems, evaluate outputs, recognize uncertainty, protect data, document assistance, and decide when artificial intelligence should not be used. These capabilities constitute responsible artificial intelligence literacy and are inseparable from disciplinary competence. Table 1 summarizes the core tensions raised in this section, setting each opportunity alongside its associated risk and the condition under which student agency is preserved.

Table 1. Core Tensions in Generative Artificial Intelligence Use in Higher Education

Opportunity	Potential Educational Value	Associated Risk	Agency Preserving Condition
Personalized explanation	Timely clarification and multiple representations	Confident but inaccurate or oversimplified understanding	Student compares explanations with course materials and explains the concept independently
Writing support	Feedback on structure, grammar, and argument	Loss of authorship, voice, and reasoning practice	Student controls claim, evidence, revisions, and final wording
Research assistance	Keyword generation and rapid orientation	Fabricated sources and shallow literature engagement	Student retrieves, reads, and verifies every cited source
Assessment support	Practice questions and formative feedback	Unauthorized completion of graded work	AI use is aligned with declared rules and followed by human demonstration
Accessibility	Translation, captioning, and flexible formats	Unequal access or unsuitable automated accommodation	Institution provides inclusive access and human alternatives
Productivity	Faster routine academic tasks	Cognitive offloading and dependence	Efficiency is used to create time for deeper analysis rather than replace it

4. Risks That Threaten Responsible Academic Agency

4.1 Epistemic Risk and the Normalization of Unverified Knowledge

Generative systems can produce fluent claims that are incomplete, misleading, or fabricated. Fluency creates an epistemic problem because students may interpret linguistic confidence as evidence of truth. This risk is especially serious when students use artificial intelligence at the beginning of research and allow generated summaries to define the boundaries of a topic. If the initial framing is distorted, later work may remain coherent but fundamentally unsound. Responsible use, therefore, requires source retrieval, comparison, and explicit uncertainty checks. Verification should be taught as a core academic practice rather than as an optional warning.

4.2 Cognitive Offloading and Skill Erosion

All learning technologies redistribute cognitive effort, and offloading is not inherently harmful. Calculators, search engines, and reference managers allow learners to redirect attention toward more complex tasks. The concern arises when students offload the very capability that an assignment is designed to develop. A student who uses artificial intelligence to correct grammar may preserve authorship. In contrast, a student who delegates problem formulation, argument development, evidence selection, and revision may bypass the entire learning process. Institutions must distinguish productive offloading from capability substitution.

4.3 Academic Integrity and Assessment Validity

Generative artificial intelligence can weaken the relationship between submitted work and student competence. When an assessment can be completed through external generation without understanding, the problem is not limited to student misconduct. It also reveals a weakness in assessment design. Universities need tasks that make thinking visible through drafts, oral explanation, applied projects, reflective commentary, local data, and iterative feedback. Authentic assessment does not mean that artificial intelligence must always be excluded. It means that the assessment should still provide credible evidence of what the student knows and can do.

4.4 Privacy, Confidentiality, and Data Governance

Students and staff may enter personal information, unpublished research, assessment materials, workplace records, or confidential institutional data into public systems without understanding retention and training practices. UNESCO guidance emphasizes the need for human-centered policy, privacy protection, and institutional validation of generative tools [23]. General-purpose risk management guidance offers a complementary resource for institutional data governance: the generative artificial intelligence profile issued by the U.S. National Institute of Standards and Technology outlines practices for identifying, measuring, and mitigating system-level risks that universities can adapt to a higher education setting [24]. Introductory guidance produced for university leaders provides a similarly practical starting point for institutions beginning to formalize policy in this area, translating high-level principles into concrete first steps for governance and staff communication [25]. Privacy literacy must therefore include practical decisions about what information may be entered, which systems are institutionally approved, and how sensitive material should be anonymized or withheld.

4.5 Bias, Cultural Representation, and Equity

Artificial intelligence systems reproduce patterns from their training data and may privilege dominant languages, institutions, and cultural assumptions. Responsible human-centered artificial intelligence requires attention to fairness, privacy, agency, transparency, and nonmaleficence [8]. These commitments align with broader ethical frameworks for artificial intelligence in society, which converge on principles of beneficence, non-maleficence, autonomy, justice, and explicability as the foundation for trustworthy deployment [26]. Earlier cross-disciplinary work establishing a unified ethical framework for artificial intelligence provides much of the conceptual grounding for these principles and for the AI4People initiative from which they emerged [27]. In higher education, this means questioning whose knowledge is represented, whether generated examples fit the local context, and whether students have equal access to high-quality tools. Equity also requires recognizing that policies written for well-resourced universities may not transfer directly to institutions with limited connectivity, infrastructure, or staff development.

4.6 Policy Ambiguity and Unequal Enforcement

Students often encounter different expectations for artificial intelligence across courses. One instructor may permit brainstorming, another may require disclosure, and another may prohibit all use. Disciplinary and assessment differences may justify variation, but unexplained inconsistency creates uncertainty and unequal enforcement. Student governance research shows support for integration together with demands for clear rules, ethical training, and meaningful consequences for misuse [28]. Universities need a common institutional vocabulary while allowing instructors to specify task-level rules.

5. The Responsible Academic Agency Framework

The Responsible Academic Agency Framework is designed to evaluate whether generative artificial intelligence strengthens or substitutes the student's role as a learner and author. The framework contains six interdependent domains. A practice should not be considered responsible merely because it satisfies one domain. For example, transparent disclosure does not compensate for unverified information, and equitable access does not guarantee that an assessment remains valid. The domains should be applied together as a design and evaluation system. The framework can operate at three points in time. Before use, it supports decision-making on purpose, permissions, data, and learning outcomes. During use, it guides verification, reflection, and human judgment. After use, it supports disclosure, explanation, evaluation, and accountability. This temporal structure encourages educators and students to treat the use of artificial intelligence as a documented academic process rather than an invisible event. Figure 1 depicts the six interdependent domains of the Responsible Academic Agency Framework and how they operate across the before, during, and after stages of artificial intelligence use.

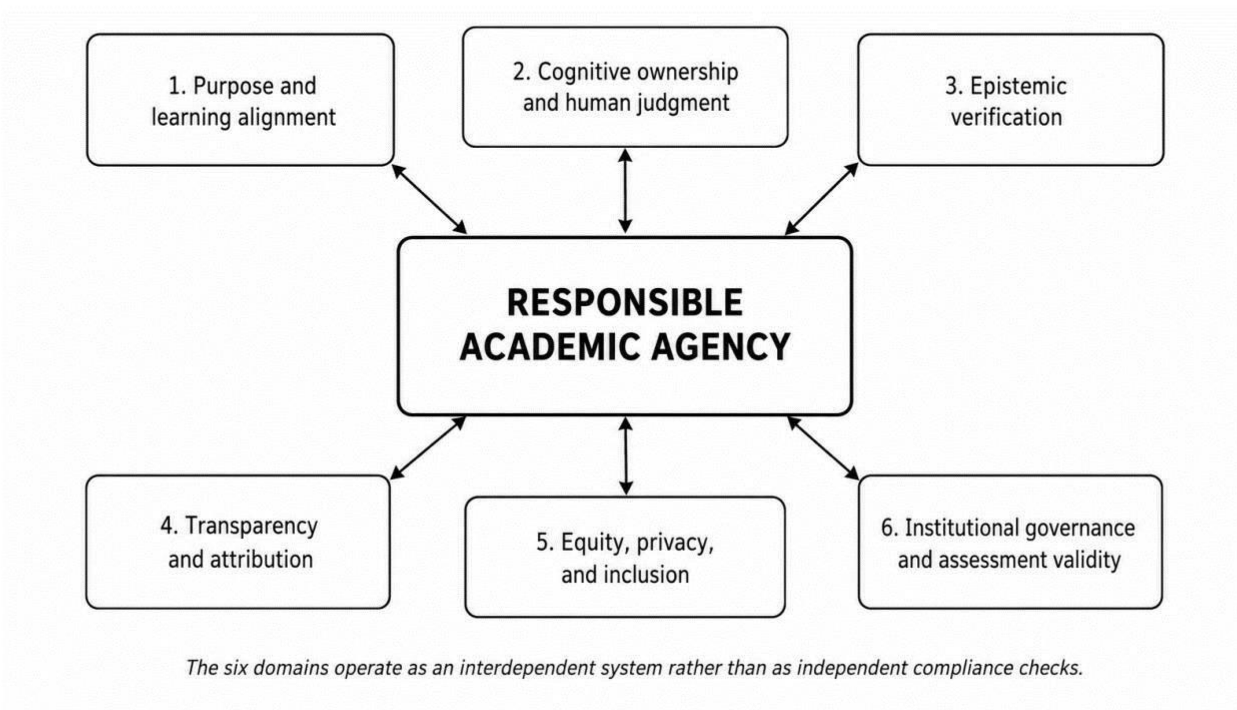

Figure 1. Responsible Academic Agency Framework for Generative AI in Higher Education

5.1 Purpose and Learning Alignment

Every use of artificial intelligence should begin with a defensible educational purpose. The student or educator should be able to explain what the tool is expected to contribute and how that contribution relates to the learning outcome. Using artificial intelligence because it is available is not a sufficient rationale. Purpose alignment asks whether the use supports understanding, practice, feedback, accessibility, or another legitimate objective. It also asks whether the tool removes the central intellectual work that the activity is intended to assess.

5.2 Cognitive Ownership and Human Judgment

Cognitive ownership means that the learner remains responsible for the direction and substance of the work. The student defines the problem, evaluates options, selects claims, interprets evidence, and decides what to accept and what to reject. Artificial intelligence may offer alternatives, but it should not become the unexamined source of purpose or conclusion. Evidence of cognitive ownership may include planning notes, prompt records, annotations, revision explanations, oral defense, and the ability to reproduce reasoning without the tool.

5.3 Epistemic Verification

Epistemic verification requires students to treat generated output as a provisional claim rather than as an authority. Factual statements should be checked against credible sources, citations should be retrieved and read, calculations should be reproduced, and uncertainty should be acknowledged. Verification is especially important when artificial intelligence is used for literature review, technical explanation, legal interpretation, health information, or data analysis. The goal is not merely to detect errors. It is to cultivate disciplined habits of evidence evaluation.

5.4 Transparency and Attribution

Transparency makes the role of artificial intelligence visible to readers, instructors, and collaborators. Disclosure should indicate the tool, the purpose of use, and the extent of material influence when required by the institution or publisher. Attribution should be proportionate. Minor spelling corrections may not require the same documentation as idea generation, structural revision, coding assistance, or text production. A clear disclosure system supports trust, but institutions should avoid treating disclosure as a substitute for evaluating learning.

5.5 Equity, Privacy, and Inclusion

Responsible use protects students from exclusion and data harm. Equity requires attention to subscription costs, device access, connectivity, disability support, language, and differential artificial intelligence literacy. Privacy requires restrictions on sensitive information and the use of institutionally reviewed tools when confidential data are involved. Inclusion requires critical attention to cultural assumptions and representational bias. These safeguards ensure that artificial intelligence integration does not reward privilege or expose vulnerable students to disproportionate risk.

5.6 Institutional Governance and Assessment Validity

Universities must create conditions in which responsible agency is possible. Governance includes clear institutional principles, course-specific expectations, staff development, approved technology pathways, data protection, student participation, and review mechanisms. Assessment validity requires evidence that the submitted work represents the student's competence.

Institutions should combine process evidence, authentic application, oral explanation, supervised performance, and reflective accounts as appropriate to the discipline. Governance should remain adaptive, as tools and practices will continue to evolve. Table 2 translates the six domains of the Responsible Academic Agency Framework into guiding questions, educational indicators, and possible evidence, offering a practical companion to recent institutional guidance for the responsible use of generative artificial intelligence [29].

Table 2. Operational Indicators for the Responsible Academic Agency Framework

Domain	**Student Question**	**Educational Indicator**	**Possible Evidence**
Purpose and learning alignment	Why am I using AI for this task?	Use has a clear relationship to the intended learning	Task rationale or permitted use statement
Cognitive ownership	Which decisions and ideas are genuinely mine?	Student controls goals, claims, interpretation, and revision	Draft history, annotations, oral explanation
Epistemic verification	How do I know this output is accurate?	Claims and sources are independently checked	Verified references, calculations, comparison notes
Transparency and attribution	What assistance must I disclose?	Material AI contribution is documented consistently	AI disclosure statement or process appendix
Equity, privacy, and inclusion	Who may be disadvantaged or exposed?	Access, confidentiality, accessibility, and cultural fit are addressed	Approved tool record, anonymization, alternative access
Governance and assessment validity	Does this work still demonstrate my competence?	Rules are clear and assessment evidence remains credible	Rubric, oral defense, practical demonstration, reflective commentary

6. Implications for Higher Education Stakeholders

6.1 Students

Students should be taught to use artificial intelligence as a critical collaborator rather than as an unquestioned authority. Practical instruction should include prompt design, source verification, error detection, privacy protection, disclosure, and reflection on cognitive dependence. Students should also be encouraged to preserve evidence of their process. This documentation can protect honest students when authorship is questioned and can help them evaluate whether artificial intelligence is improving or replacing their learning.

6.2 Educators

Educators need support to define acceptable use at the level of specific tasks. General statements such as use artificial intelligence responsibly are too vague. Assignment instructions should state what is permitted, what requires disclosure, what is prohibited, and how understanding will be demonstrated. Educators should design activities that reward judgment, local application, engagement with sources, and revision. They should also disclose significant use of artificial intelligence in their teaching materials to model the transparency expected from students.

Faculty readiness for this role is uneven. Research on instructors' use and self-efficacy shows that professional development needs vary considerably across disciplines and career stages, so institutions cannot assume that educators are equally prepared to model responsible practice without targeted support [30].

6.3 University Leaders and Quality Assurance Units

Institutional leaders should move from emergency policy toward coherent governance. A university-wide policy should establish principles and terminology, while faculties and instructors translate those principles into disciplinary practice. Quality assurance processes should examine whether programs develop artificial intelligence literacy, whether assessments remain valid, whether approved systems protect data, and whether students have equitable access. Students should participate in policy development because they experience the practical consequences of ambiguity and enforcement. Institutional frameworks developed specifically for responsible generative artificial intelligence use in research contexts offer a useful complement to the teaching-focused governance discussed here, suggesting that leaders may benefit from aligning teaching, research, and administrative policies under shared principles rather than treating them as separate initiatives.

6.4 Libraries, Writing Centers, and Academic Support Services

Libraries and academic support units are central to responsible implementation. Librarians can teach source verification, citation tracing, and information evaluation. Writing centers can help students distinguish between feedback and substitution and preserve their individual voice. Academic support services can also provide alternatives for students who cannot access paid systems or who require disability accommodations. These units should be integrated into an artificial intelligence strategy rather than treated as peripheral services.

6.5 Technology Providers

Technology providers should support educational accountability by implementing privacy controls, ensuring source visibility, communicating uncertainty, ensuring accessibility, and configuring institutions. Systems designed for education should make it easier to trace evidence, preserve student work, and understand model limitations. Product design should not assume that maximum convenience is always educationally desirable. In some learning contexts, friction, questioning, and delayed assistance may better support agency than immediate answer generation.

7. Research Agenda

The next phase of research should move beyond descriptive studies of attitudes and usage frequency. Adoption data are important, but they do not show whether students are learning more deeply, verifying information, retaining skills, or becoming more capable of independent judgment. Research should examine how different patterns of artificial intelligence use influence academic agency over time. Longitudinal designs are especially necessary. Short studies may identify immediate efficiency or satisfaction, but miss later dependence or skill change. Researchers should compare students who use artificial intelligence for explanation, feedback, generation, and verification, rather than treating all uses as one category. Studies should also include disciplinary differences because acceptable cognitive offloading in one field may be central competence in another. Greater attention is needed in the Global South and in institutions with limited resources. Much policy and empirical research remains concentrated in highly resourced contexts, even though access, language, infrastructure, and governance conditions vary substantially. Comparative work should avoid assuming that a single institutional model is universally appropriate. It should examine how a responsible academic agency can be supported under different technological and cultural conditions. Table 3 organizes this research agenda into priority areas, illustrative questions, and the kinds of evidence that would advance each one.

Table 3. Priority Research Questions

Research Area	**Illustrative Question**	**Suggested Evidence**
Learning and cognition	Which AI use patterns strengthen or weaken independent reasoning and long-term retention?	Longitudinal performance, think aloud protocols, delayed tests
Verification behavior	What instructional designs increase students' detection and correction of AI errors?	Source tracing tasks, error identification, process logs
Assessment validity	Which assessment combinations provide credible evidence of competence when AI is available?	Oral defense, authentic projects, supervised tasks, rubric studies
Equity	How do cost, infrastructure, language, disability, and prior literacy shape AI benefits and risks?	Cross-institutional and cross national comparisons
Policy communication	How do students interpret inconsistent institutional and course level AI rules?	Policy audits, interviews, scenario experiments
Agency development	Can responsible AI instruction increase? student autonomy rather than dependence?	Validated agency scales, self-regulation measures, qualitative trajectories
Faculty practice	How does educator AI use influence? student trust and perceptions of fairness?	Disclosure experiments, classroom studies, trust measures
Governance	Which institutional arrangements adapt effectively as AI systems change?	Longitudinal case studies, governance maturity indicators

8. DISCUSSION

The Responsible Academic Agency Framework changes the unit of analysis in higher education artificial intelligence research. Instead of asking only whether a student used a tool, it asks what intellectual and ethical responsibilities remained with the student. This distinction is essential because the same observable output may arise from very different learning processes. A polished essay may represent careful student reasoning supported by language feedback, or it may represent extensive generation accepted without understanding. Product inspection alone cannot reliably distinguish these pathways. The framework also avoids a false choice between innovation and integrity. Artificial intelligence can support learning while academic standards remain rigorous, but only when institutions align tool use with purpose, verification, transparency, equity, and valid assessment. Conversely, prohibition does not automatically protect the agency. Students may use tools secretly, policies may be enforced inconsistently, and conventional assignments may continue to reward reproduction rather than understanding. The framework, therefore, treats responsible integration as a design problem rather than a surveillance problem. A further contribution is the integration of individual and institutional responsibility. Students remain accountable for their work, yet institutions shape the conditions under which responsible decisions are possible. Ambiguous rules, inaccessible support, weak privacy controls, and invalid assessments cannot be solved through student ethics alone. Responsibility must be distributed across the sociotechnical system while preserving clear human accountability.

The framework has limitations. It is conceptual and requires empirical validation across disciplines, institutional types, and cultural settings. Its domains may also need adaptation as artificial intelligence becomes more agentic and capable of acting across multiple academic systems. Nevertheless, the focus on purpose, judgment, evidence, transparency, equity, and governance provides a stable foundation, as these concerns arise from the educational mission rather than the temporary capabilities of a particular model.

9. CONCLUSION

Generative artificial intelligence is not merely changing how students complete academic work. It is changing where thinking occurs, how authorship is demonstrated, and how universities judge competence. The appropriate response is neither unconditional adoption nor universal prohibition. Higher education needs a principled approach that preserves students' roles as intentional, reflective, and accountable learners. A responsible academic agency offers such an approach. It requires artificial intelligence use to be aligned with learning purposes, governed by human judgment, grounded in verified evidence, disclosed transparently, implemented equitably, and supported by valid assessment and institutional accountability. Under these conditions, artificial intelligence can expand access to feedback and intellectual exploration without becoming a substitute for learning. The future credibility of higher education will depend not on whether universities use artificial intelligence, but on whether graduates can still demonstrate the knowledge, judgment, integrity, and responsibility that university education is intended to develop.

DECLARATIONS

CONFLICT OF INTEREST: The author declares no known financial or personal conflict of interest.

FUNDING: No external funding was received for this conceptual article.

ETHICS APPROVAL: Ethics approval was not required because the article did not involve human participants or identifiable personal data.

DATA AVAILABILITY: No original dataset was generated or analyzed for this conceptual article.

ARTIFICIALINTELLIGENCE DISCLOSURE: Generative artificial intelligence was used as a language and organizational support tool during manuscript development. The author remains responsible for the conceptual argument, source verification, interpretation, and final text. This disclosure should be revised to match the policy of the selected journal before submission

10. REFERENCES

[1] Bond M, Khosravi H, De Laat M, Bergdahl N, Negrea V, Oxley E, et al. A meta-systematic review of artificial intelligence in higher education: a call for increased ethics, collaboration, and rigor. Int J Educ Technol High Educ. 2024;21:4. doi:10.1186/s41239-023-00436-z.

[2] Zawacki-Richter O, Marín VI, Bond M, Gouverneur F. Systematic review of research on artificial intelligence applications in higher education: where are the educators? Int J Educ Technol High Educ. 2019;16:39. doi:10.1186/s41239-019-0171-0.

[3] Crompton H, Burke D. Artificial intelligence in higher education: the state of the field. Int J Educ Technol High Educ. 2023;20:22. doi:10.1186/s41239-023-00392-8.

[4] Al-Zahrani AM, Alasmari TM. Exploring the impact of artificial intelligence on higher education: the dynamics of ethical, social, and educational implications. Humanit Soc Sci Commun. 2024;11:912. doi:10.1057/s41599-024-03432-4.

[5] Slimi Z. The impact of artificial intelligence on higher education: an empirical study. Eur J Educ Sci. 2023;10(1):17-33.

[6] Bittle K, El-Gayar O. Generative AI and academic integrity in higher education: A systematic review and research agenda. Information.2025;16(4):296. doi:10.3390/info16040296.

[7] Francis NJ, Jones S, Smith DP. Generative AI in higher education: balancing innovation and integrity. Br J Biomed Sci. 2025;82:14048. doi:10.3389/bjbs.2024.14048.

[8] Fu Y, Weng Z. Navigating the ethical terrain of AI in education: a systematic review on framing responsible human-centered AI practices. Comput Educ Artif Intell. 2024;7:100306. doi:10.1016/j.caeai.2024.100306.

[9] Kasneci E, Sessler K, Küchemann S, Bannert M, Dementieva D, Fischer F, et al. ChatGPT for good? On opportunities and challenges of large language models for education. Learn Individ Differ.2023;103:102274. doi:10.1016/j.lindif.2023.102274.

[10] Zhu H, Sun Y, Yang J. Towards responsible artificial intelligence in education: a systematic review on identifying and mitigating ethical

risks. Humanit Soc Sci Commun. 2025;12:1001. doi:10.1057/s41599-025-05252-6.

[11] Chan CKY. A comprehensive AI policy education framework for university teaching and learning. Int J Educ Technol High Educ. 2023;20:38. doi:10.1186/s41239-023-00408-3.

[12] Jin Y, Yan L, Echeverría V, Gašević D, Martínez-Maldonado R. Generative AI in higher education: a global perspective of institutional adoption policies and guidelines. Comput Educ ArtifIntell.2025;8:100348. doi:10.1016/j.caeai.2024.100348.

[13] Escotet MA. The optimistic future of artificial intelligence in higher education. Prospects. 2024;54(3):531-40. doi:10.1007/s11125-023-09642-z.

[14] Bandura A. Social cognitive theory: An agentic perspective. Annu Rev Psychol. 2001;52:1-26. doi:10.1146/annurev.psych.52.1.1.

[15] Zimmerman BJ. Becoming a self-regulated learner: an overview. Theory Pract. 2002;41(2):64-70. doi:10.1207/s15430421tip4102_2.

[16] Long DY, Wang S, Md Rashid S, Lu XT. Artificial intelligence in higher education: a systematic review of its impact on student engagement and the mediating role of teaching methods. Front Educ. 2026;10:1648661. doi:10.3389/feduc.2025.1648661

[17] Black RW, Tomlinson B. University students describe how they adopt AI for writing and research in a general education course. Sci Rep. 2025;15:8799. doi:10.1038/s41598-025-92937-2.

[18] Weaver KD. The artificial intelligence disclosure framework: an introduction. Coll ResLibrNews.2024;85(4):150-3. doi:10.5860/crln.85.4.150.

[19] Cotton DRE, Cotton PA, Shipway JR. Chatting and cheating: ensuring academic integrity in the era of ChatGPT. Innov Educ Teach Int. 2024;61(2):228-39. doi:10.1080/14703297.2023.2190148.

[20] Jafari F, Keykha A. Identifying the opportunities and challenges of artificial intelligence in higher education: A qualitative study. J Appl Res High Educ.2024;16(4):1228-45.doi:10.1108/JARHE-09-2023-0426.

[21] Nikolopoulou K. Generative artificial intelligence in higher education: exploring ways of harnessing pedagogical practices with the assistance of ChatGPT. Int J Chang Educ. 2024;1(2):103-11. doi:10.47852/bonviewIJCE42022489.

[22] Tlili A, Shehata B, Adarkwah MA, Bozkurt A, Hickey DT, Huang R, et al. What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education? Smart Learn Environ. 2023;10:15. doi:10.1186/s40561-023-00237-x.

[23] Miao F, Holmes W. Guidance for generative AI in education and research. Paris: UNESCO; 2023.

[24] National Institute of Standards and Technology. Artificial intelligence risk management framework: generative artificial intelligence profile (NIST AI 600-1). Washington (DC): U.S. Department of Commerce; 2024. doi:10.6028/NIST.AI.600-1.

[25] Sabzalieva E, Valentini A. ChatGPT and artificial intelligence in higher education: quick

start guide. Caracas: UNESCO International Institute for Higher Education in Latin America and the Caribbean; 2023.

[26] Floridi L, Cowls J. A unified framework of five principles for AI in society. Harv Data Sci Rev. 2019;1(1). doi:10.1162/99608f92.8cd550d1.

[27] Floridi L, Cowls J, Beltrametti M, Chatila R, Chazerand P, Dignum V, et al. AI4People: an ethical framework for a good AI society. Minds Mach. 2018;28:689-707. doi:10.1007/s11023-018-9482-5

[28] Barus OP, Hidayanto AN, Handri EY, Sensuse DI, Yaiprasert C. Shaping generative AI governance in higher education: insights from student perception. Int J Educ Res Open. 2025;8:100452. doi:10.1016/j.ijedro.2025.100452.

[29] Smith SM, Tate M, Freeman K, Walsh A, Ballsun-Stanton B, Lane M. A university framework for the responsible use of generative AI in research. J High Educ Policy Manag. 2026;48(1):17-36. doi:10.1080/1360080X.2025.2509187.

[30] Mah DK, Groß N. Artificial intelligence in higher education: exploring faculty use, self-efficacy, distinct profiles, and professional development needs. Int J Educ Technol High Educ. 2024;21:58. doi:10.1186/s41239-024-00490-1.