



# The Tension Between AI-Driven Knowledge Building and Critical Thinking among College Students in Tanzania

Leonard Nyanzira

Department of Management Studies

Tanzania Institute of Accountancy (TIA), Mwanza Campus, Tanzania

Corresponding email: [leonardnyanzira@gmail.com](mailto:leonardnyanzira@gmail.com)

ORCID ID - <https://orcid.org/0000-0002-7247-8893>

**Abstract:** This study investigates the dual role of generative AI (GAI) tools in higher education, focusing on the tension between knowledge construction and critical thinking. A survey was administered to students in higher learning institutions in Mwanza and Dodoma, Tanzania. Quantitative findings reveal that GAI is deeply integrated into learning routines, with a majority of respondents using it daily or weekly primarily for research, problem-solving, and simplifying complex concepts. However, a significant concern emerged, with over 60% of students identifying "reduced critical thinking" as a primary risk, highlighting an acute awareness of the potential for passive consumption and over-reliance. Crucially, students are proactively developing metacognitive strategies to mitigate these risks, with the most common being verification of AI outputs against credible sources and critical prompting. The study concludes that while GAI serves as a powerful scaffold for efficient knowledge construction, its perceived threat to critical thinking is substantial. These findings underscore the urgent need for pedagogical frameworks that integrate "critical AI literacy" to help students harness these tools effectively without compromising higher-order cognitive skills.

**Keywords:** Generative AI, College students, knowledge building, critical thinking

## How to cite this work (APA):

Nyanzira, L. (2026). The Tension Between AI-Driven Knowledge Building and Critical Thinking among College Students in Tanzania. *Journal of Research Innovation and Implications in Education*, 10(1), 58 – 67. <https://doi.org/10.59765/lp64w>.

## 1. Introduction

The adoption of Generative Artificial Intelligence (GAI) into higher education has sparked a global debate about its potential and inherent risks. Tools such as ChatGPT, Gemini, and discipline-specific copilots now sit alongside textbooks and lecture notes as everyday study companions. Early syntheses describe GAI as a "reading/writing partner" that compresses the time to understanding, scaffolds complex language, and supports ideation across tasks ranging from outlining to code explanation (Kasneci et al., 2023). Students report immediate benefits, clearer explanations, faster progress, and reduced anxiety, which aligns with classic Uses &

Gratifications theory: learners select media that satisfy informational and efficiency needs in context (Ruggiero, 2000; Sundar & Limperos, 2013). At the same time, a counter-current of work warns that these efficiencies may come with cognitive costs, especially if GAI substitutes for careful reading, argument construction, and self-monitoring, which is the core of critical thinking (Barshay, 2025; Pohl et al., 2024).

This tension is especially salient in African higher education, where students often study under infrastructural constraints, such as irregular connectivity, uneven library holdings, and high data costs. Regional research shows that learners creatively "patch" these gaps by using digital tools to translate, condense, and

reformat material to fit local needs (Gondwe & Awami, 2025; Mare, 2015). In such settings, the attraction of GAI is not merely convenience; it is a pragmatic response to material limits. That response, however, can also intensify the risk of over-reliance: when time, power, and bandwidth are scarce, polished AI outputs can be accepted too quickly, with fewer opportunities for deep reading or triangulation. Platform studies further suggest that the fluent style of GAI (such as confident, well-formed prose and step-by-step reasoning) can confer “algorithmic authority,” encouraging uncritical uptake (Gondwe, 2025; Gillespie, 2014) and amplifying known dynamics of automation bias, in which users overweight system outputs, particularly under time pressure (Parasuraman & Riley, 1997). We conceptualize the core problem as a trade-off between AI-driven knowledge building, (the rapid formation of baseline understanding via search, summarization, brainstorming, translation, and step-by-step guidance), and critical thinking, defined here as analysis, evaluation of evidence, synthesis across sources, and reflective judgment. Cognitive offloading theory suggests that people routinely shift memory and procedure to external artifacts (Risko & Gilbert, 2016). Offloading can be productive, freeing capacity for higher-order work, but chronic dependence can narrow practice opportunities and dull metacognitive calibration. Studies of online reasoning add a second risk: many learners already struggle with source evaluation and lateral reading, skills that become even more important when working with fluent but fallible systems (Wineburg & McGrew, 2017). In politically contested information environments, as in parts of Southern Africa, the stakes of gullibility are higher because of histories of surveillance and propaganda condition both skepticism and self-censorship (Chitanana, 2024).

Yet the emerging literature also documents a constructive counter-trend. Students and instructors are articulating critical AI literacy through a set of metacognitive practices that treat GAI as a starting point rather than an answer. As such, critical prompting, explicit verification against trusted sources, and reflective use that keeps the human “in the loop” (URI ITS, 2024). When paired with self-assessment and goal-setting, these practices have been linked to deeper learning (Fischer et al., 2018). Importantly, disciplinary context matters. Procedural and correctness-oriented fields (e.g., mathematics, computing, engineering) often emphasize accuracy and step validity, making error-propagation and over-reliance salient concerns; by contrast, fields that privilege synthesis and originality (e.g., law, social sciences, humanities) often center the risks of diminished independent thought and homogenized voice (Pohl et al., 2024; Hechinger Report, 2025). Despite fast-growing global interest, three gaps limit what we know. First, most empirical studies come from North America and Europe, with far fewer accounts from Sub-Saharan Africa. This geographic skew risks universalizing experiences from high-resource settings and overlooking how connectivity, cost, and local

pedagogies shape the calculus of AI use (Mare, 2015; Gondwe & Awami, 2025). Second, while researchers advocate for critical AI literacy, few studies systematically document the bottom-up strategies students already employ in African contexts, how verification, prompt design, and resource mixing are enacted in practice. Third, disciplinary comparisons outside Western settings are rare, leaving open the question of whether observed STEM–Humanities differences generalize across infrastructure and culture.

This article addresses these gaps with evidence from higher education institutions in Mwanza and Dodoma, Tanzania. Using a cross-sectional survey with both closed and open items, we examine how students integrate GAI into routine study, what they see as the main risks, which metacognitive practices they report using to stay critically engaged, and how these patterns vary by disciplinary domain. Our approach treats students as active agents rather than passive adopters: we ask not only *what* they do with GAI, but *how* they try to keep their own judgment in the loop. Our contributions are threefold. Conceptually, we refine the efficiency–rigor trade-off by situating it within infrastructural realities. We argue that the same behavior (e.g., adopting AI summaries) can function as *adaptive patching* in constrained settings even as it introduces risks of superficiality, placing design and policy responsibilities on institutions rather than individual students alone. Empirically, we provide a granular accounting of student-initiated critical AI literacy practices (e.g., verification routines, “tutor-not-shortcut” framing, purposeful prompting), extending prior work that has largely theorized these strategies without local detail (URI ITS, 2024; Fischer et al., 2018). Comparatively, we test whether disciplinary risk profiles observed in Western studies appear in a Tanzanian sample, offering evidence to inform discipline-sensitive pedagogy (Pohl et al., 2024; Hechinger Report, 2025). Guided by these aims, the study addresses three research questions relating to the extent the use of generative AI for knowledge construction relates to students’ perceptions of reduced critical thinking; the metacognitive strategies students report using to mitigate over-reliance and ensure critical engagement with AI-generated content, and the perceived utilities and risks of GAI differentials between quantitative/STEM and qualitative/humanities students. By centering Tanzanian students’ experiences, we seek to move beyond binary debates over whether GAI is “good” or “bad” for learning. Instead, we map the concrete trade-offs students perceive, the safeguards they already practice, and the discipline-specific contours of risk. This reframing has direct implications for curriculum, assessment, and policy. Rather than prohibition or uncritical embrace, our findings argue for integrating critical AI literacy as a core academic skill—taught explicitly, assessed transparently, and adapted to local infrastructure and disciplinary epistemologies (Gillespie, 2014; Risko & Gilbert, 2016; Wineburg & McGrew, 2017; URI ITS, 2024). In doing so, we

contribute evidence toward a pragmatic middle path: harness GAI as a powerful scaffold for knowledge building while deliberately protecting and cultivating the higher-order thinking that higher education exists to develop.

## 1.2 Research Objectives

Reflecting on the main purpose mentioned above, this study focuses on the following specific objectives.

**RO1:** To determine the extent to which the use of generative AI for knowledge construction influences the development of critical thinking skills among higher education students.

**RO2:** To identify and examine the metacognitive strategies that students use to minimize over-reliance on generative AI and to maintain critical engagement when interacting with AI-generated content.

**RO3:** To compare perceptions of the benefits and risks of generative AI use between students in quantitative/STEM disciplines and those in qualitative/humanities disciplines.

## 1.3 Research Questions

Against the background, we ask the following research questions:

**RQ1:** To what extent does the use of generative AI for knowledge construction impact the development of critical thinking skills in higher education students?

**RQ2:** What metacognitive strategies do students employ to mitigate the risks of over-reliance and ensure critical engagement with AI-generated content?

**RQ3:** How do perceptions of AI's utility and risks differ between students in quantitative/STEM fields and those in qualitative/humanities disciplines.

## 2. Literature Review

The advent of generative artificial intelligence (GAI) tools has made a significant paradigm shift in higher education globally. Emerging studies continue to suggest that while these technologies offer unparalleled opportunities for democratizing information and accelerating academic tasks, they also introduce a tension between the potential to enhance "knowledge acquisition and the cost of critical thinking (Stuart, 2024; Sarakikya & Kitula, 2024; Mambile & Mwagosi, 2025). Early syntheses in higher education portray GAI as a "writing/reading partner" that compresses time to understanding, scaffolds complex language, and supports ideation (e.g., Kasneci et al., 2023). Studies across disciplines show strong uptake for information finding, drafting, summarizing, and problem

decomposition; STEM learners commonly request worked steps or code explanations, while humanities students emphasize outlining, paraphrase, and argument templates. Through these uses, students report immediate benefits of clarity, speed, reduced anxiety, and a sense of progress on otherwise sticky tasks. This trend echoes the classic *Uses & Gratifications* logic in which learners select media that satisfy informational and efficient needs (Ruggiero, 2000; Sundar & Limperos, 2013).

In African higher education, AI's appeal is amplified by longstanding frictions in access to scholarly materials, uneven library holdings, and high transaction costs for finding quality sources. Regional work shows students taking advantage of digital tools to "patch" infrastructural gaps by translating, condensing, and reformatting content for local needs (Gondwe & Awami, 2025). Comparative studies in Southern Africa also highlight how platforms intertwine with cost, connectivity, and device sharing, shaping which tasks are offloaded into technology and which remain human-performed (Mare, 2015). Regardless of the growing body of academic literature on AI in education in Africa, most of it is heavily skewed toward Western contexts rooted in three thematic areas, namely the dual nature of GAI, the rise of student-initiated metacognitive strategies, and the disciplinary differences in AI use. Research consistently highlights the efficiency gains GAI provides for knowledge construction. Studies have shown that students use tools like ChatGPT for tasks ranging from quick information retrieval and problem-solving to brainstorming and summarization (Rahman & Watanobe, 2023; Sánchez-Ruiz et al., 2023). The ability of these tools to generate well-structured and contextually relevant responses on demand has been found to significantly improve the speed and quality of academic work, acting as a powerful scaffold for learners (Al-Dhaqri et al., 2024). However, a robust and growing body of literature argues that this efficiency comes at a significant cognitive cost. The fear of "cognitive offloading" is a central theme (Barshay, 2025). Essentially, Barshay argues that when students consistently rely on GAI to perform mental tasks (such as synthesizing disparate information or structuring complex arguments) they may fail to develop these skills themselves.

A study by The Hechinger Report (2025) also found that students with AI access were less likely to engage directly with primary reading materials, instead opting to query the chatbot, an observation that points to the outsourcing of critical text analysis. For example, a European study found that 83% of students expressed concern that excessive reliance on AI could weaken their ability to think independently (Pohl et al., 2024).<sup>9</sup> These findings echo the concerns raised in the user's in this study, demonstrating that the tension between knowledge building and critical thinking is a widely recognized phenomenon. While much of the early discourse focused on the risks, a more recent and

promising body of literature is shifting focus to how students are actively adapting to these challenges. The concept of "critical AI literacy" has emerged, emphasizing the importance of self-regulatory practices such as "critical prompting" and "source verification" (URI ITS, 2024). Instead of accepting GAI outputs as final answers, students are being advised to treat them as a starting point for inquiry, cross-referencing claims with peer-reviewed articles, textbooks, and other reliable sources. This approach transforms GAI from a passive "answer machine" into an active "tutor" or "study companion," where the student remains in the driver's seat. Researchers have shown that pairing AI with explicit metacognitive strategies, such as self-assessment and goal-setting, can significantly improve learning outcomes and promote a deeper understanding (Fischer et al., 2018).<sup>11</sup> The user's study's findings on verification practices and the "tutor, not shortcut" mindset directly align with this emerging body of literature, suggesting that these behaviors are not isolated incidents but part of a broader, student-led response to the challenges of GAI.

Recent scholarship has also begun to disaggregate AI's impact across different academic disciplines. The needs and challenges of a STEM student working on a coding problem are fundamentally different from those of a humanities student writing a critical essay. A study by The Hechinger Report (2025) noted that students in disciplines requiring factual or procedural answers (like mathematics or engineering) might be more concerned with the accuracy and over-reliance on the AI's steps. In contrast, those in qualitative fields (like law or social sciences) might worry more about originality and the loss of independent thought. This is because the core cognitive demands of these fields differ: one prioritizes correctness and procedural logic, while the other values synthesis, argumentation, and originality (Pohl et al., 2024). Nonetheless, three theoretical strands of research converge on the concern that efficiency may crowd out critical thought. These include automation bias and algorithmic authority, cognitive offloading, and epistemic cognition and misinformation.

In automation bias and algorithmic authority, human-automation interaction research shows that when systems deliver confident, fluent answers, users overweight those outputs (especially under time pressure), substituting system judgment for their own (Parasuraman & Riley, 1997). Platform studies add that the *format* of GAI responses (polished prose, coherent reasoning steps) confers "algorithmic authority," inviting acceptance without verification (Gillespie, 2014). On the other hand, work in cognitive psychology demonstrates that people strategically offload memory and procedural steps to external artifacts (Risko & Gilbert, 2016). Offloading can be beneficial, but chronic reliance reduces opportunities for practice and metacognitive calibration. In learning contexts, repeated dependence on solver-style outputs can erode students' ability to monitor reasoning quality, generate counter-arguments, or

transfer knowledge to novel contexts. Further, in regard to **epistemic cognition and misinformation**, studies of online reasoning show that many learners struggle with source evaluation and lateral reading, skills that are vital when interacting with fluent but fallible systems (Wineburg & McGrew, 2017). Where information ecologies are politically contested, as in parts of Southern Africa, the stakes of gullibility are higher; students' self-censorship and skepticism are shaped by local histories of surveillance, propaganda, and platform governance (Chitanana, 2024). Together, these theories specify *how* and *why* GAI's strengths (speed, fluency, task completion) can dampen the very practices, thus questioning, evidencing, and weighing alternatives, that constitute critical thinking.

## 2.1 Research Gaps

Despite the growing body of literature, several critical gaps remain, particularly regarding the experiences of non-Western student populations. These gaps include the lack of contextualized research in Sub-Saharan Africa, under-researched metacognitive strategies, and comparisons of disciplinary differences in a non-Western context. This study is deliberately designed to address these lacunae, making several direct contributions to the academic discourse. The vast majority of empirical studies on GAI in higher education have been conducted in North America, Europe, and parts of Asia. This leaves a significant scarcity of research exploring the unique challenges and opportunities of GAI in African contexts, such as Tanzania. Issues such as limited access to reliable internet, varying levels of digital literacy, and diverse pedagogical traditions are often overlooked. This study is a crucial first step in filling this void, providing rich, localized data from Mwanza and Dodoma that can inform educational policy and pedagogical approaches tailored to the specific needs and infrastructure of Tanzania and other Sub-Saharan African nations.

Second, while the literature acknowledges the importance of metacognition, few studies have empirically identified and categorized the specific, on-the-ground strategies that students are spontaneously developing to cope with GAI's challenges. The user's study's detailed findings on "verification practices" and the "tutor, not shortcut" mindset are a direct contribution, moving beyond a theoretical understanding to provide concrete examples of student agency. This is a critical departure from the often-pessimistic "AI-as-threat" narrative. Third, the user's study provides a valuable, data-driven comparison of how students in quantitative/STEM and qualitative/humanities fields perceive and use GAI. This adds a new layer to the existing literature by confirming that disciplinary differences in AI perception are not just a Western phenomenon but are likely rooted in universal disciplinary demands. This study, therefore, challenges the assumption that AI's impact is monolithic and instead

calls for more localized and discipline-specific pedagogical responses.

## **3. Methodology**

### **3.1 Research Design**

This study employed a cross-sectional survey design to investigate how higher education students in Tanzania use generative artificial intelligence (AI) tools, with a particular focus on the relationship between rapid knowledge construction and critical thinking. A survey approach was chosen because it allowed for the efficient collection of large-scale self-reported data on attitudes, behaviors, and strategies across multiple disciplines.

### **3.2 Samples and Sampling**

The target population consisted of undergraduate and graduate students enrolled in higher education institutions located in Mwanza and Dodoma, two of Tanzania's major urban centers with diverse academic offerings. An estimated 1,000 students were invited to take part in the study through institutional mailing lists, academic WhatsApp groups, and direct outreach by faculty collaborators. Out of these, 200 students completed the online survey, representing a response rate of roughly 20%. After screening for eligibility and data quality, including removal of incomplete responses, patterned answers, or responses from individuals not currently enrolled in higher education, which led to a final sample of 156 students retained for analysis.

The demographic distribution of the sample covered both genders, a range of academic years, and a mix of disciplinary fields, including quantitative programs such as accounting, finance, ICT, and mathematics, as well as qualitative programs such as journalism, linguistics, education, and law. This diversity enabled comparisons across disciplinary domains.

### **3.3 Data Collection Tools**

The study employed an online questionnaire consisting of both closed-ended and open-ended items. The survey included items organized around three conceptual domains:

#### **3.3.1 Knowledge Construction**

Questions captured how students used AI for learning, including research and information gathering, summarization, translation, brainstorming, problem solving, and exam preparation.

#### **3.3.2 Critical Thinking Concerns**

Respondents were asked to select or describe perceived risks associated with AI use, such as misinformation, over-reliance, reduced critical thinking, internet dependency, and lack of personalization.

#### **3.3.3 Metacognitive Strategies**

Students were asked about self-regulatory practices, such as verifying information with textbooks, setting learning goals, asking specific questions, or using AI as a tutor rather than a shortcut.

To allow for disciplinary comparisons, students also reported their field of study, which was coded into broad categories of quantitative/STEM or qualitative/humanities disciplines.

### **3.4 Data Collection Procedure**

The survey was administered online between May and September 2025 using a secure web-based platform. Participants received an invitation link containing information about the purpose of the study, confidentiality assurances, and voluntary participation guidelines. Informed consent was obtained electronically before participants could proceed to the questionnaire. The survey instrument consisted of both closed-ended and open-ended questions. Closed-ended questions asked students to report on their frequency of AI use, the tools they relied on most frequently (e.g., ChatGPT, Google Bard/Gemini, DeepSeek), and their perceived benefits and risks (e.g., misinformation, reduced critical thinking, internet dependency). Open-ended questions asked students to describe in their own words how they used AI for academic purposes, the strategies they employed to maintain critical engagement, and any challenges they experienced.

### **3.5 Data Analysis**

Quantitative data was analyzed using descriptive statistics (frequency counts, percentages, and confidence intervals) to map the prevalence of different uses and concerns. Cross-tabulations were performed to compare responses between quantitative/STEM and qualitative/humanities students. Qualitative data from open responses were thematically coded to identify recurring strategies and explanatory statements that contextualized the numerical patterns. The integration of quantitative and qualitative evidence provided a richer understanding of how students balanced knowledge construction with critical thinking in their use of generative AI.

### **3.6 Ethical Considerations**

The study adhered to standard ethical guidelines for educational research. Participation was voluntary, and respondents provided informed consent electronically.

No identifying information was collected, ensuring anonymity and confidentiality. The data were stored securely and used solely for academic research purposes. Participants were informed of their right to withdraw at any point before submitting the survey. The research complied with institutional and national ethical requirements for studies involving human participants.

## 4. Results and Discussion

This study provides a detailed look at the use of Generative AI (GAI) among college students in Tanzania. The survey shows that generative AI has moved from the margins into everyday academic work, helping students build knowledge quickly while also raising a steady worry that it can dull the habits of analysis and judgment. Drawing on the full set of responses addressing our research questions, we first describe the overall pattern of use and perceived risks, then turn to what students themselves do to stay critically engaged and finally compare how these views differ between quantitative/STEM programs and qualitative/humanities fields.

### RQ1. Knowledge construction and its effect on critical thinking

This question set out to determine how students' use of generative AI relates to their own sense that their critical thinking may be weakened or displaced. In other words, the aim was to see whether speeding up the "knowledge construction" part of learning comes with a perceived cost to analysis, reflection, and judgment. Across the dataset, students overwhelmingly describe AI as a fast scaffold for learning. They most often use it to search for information, solve problems, generate ideas, summarize difficult texts, and translate difficult language. Mentions of these activities are frequent: research and information gathering appear in about three-quarters of responses, solving academic problems in roughly two-thirds, brainstorming in a little under three-fifths, and summarizing and translation in just under half. Many students say that these uses help them grasp the basics more quickly, organize their thinking, and move past sticking points in assignments. Figure 1 summarizes the findings:

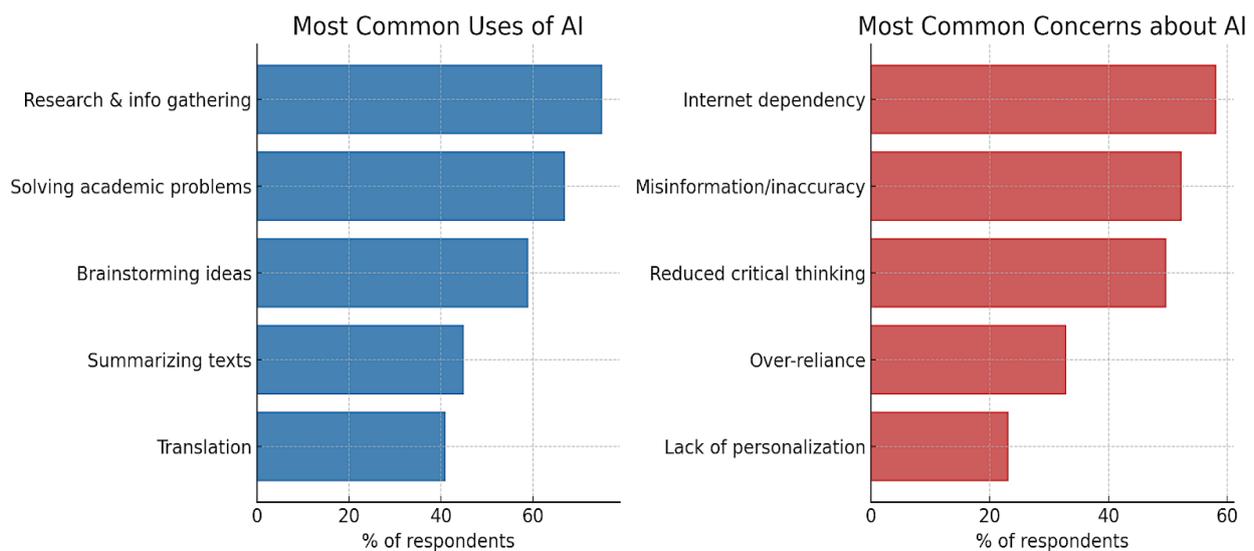


Figure 1: Most common concerns about AI in Higher education

This graph presents a side-by-side design that makes it easy to see the central tension: AI is valued for speed and convenience, but those very strengths are tied to risks that students themselves notice. On the left, it shows the most common ways students use AI, with research, problem-solving, and brainstorming being the top activities. On the right, it shows the major concerns, with internet dependency, misinformation, and reduced critical thinking dominating the picture.

At the same time, nearly one in two students explicitly names "reduced critical thinking" as a risk of using AI in this way. In numerical terms, 49.7% of respondents mention this concern (95% confidence interval 41.8% to

57.5%). Other concerns travel with it: 58.1% mention internet dependence, 52.3% mention misinformation or inaccuracy, 32.9% mention over-reliance, and 23.2% mention lack of personalization. Reading the open responses together with these counts, a common pattern emerges. Students value AI for compressing the time it takes to form a baseline understanding, but they also recognize that this very efficiency can invite a "shortcut mindset" in which they accept polished answers too quickly. The language they use (such as "use as a study companion, not a shortcut," "guidance not answers," and "don't rely 100%") shows that many are alert to the risk of outsourcing judgment even as they benefit from AI's speed. ChatGPT is the dominant tool by mention (about

nine in ten responses), with Google Bard/Gemini and DeepSeek appearing less often. This distribution does not change the core pattern that the quicker the initial understanding, the more students feel they must actively protect their critical engagement.

### **RQ2. Metacognitive strategies that limit over-reliance and sustain critical engagement**

RQ2 sought to identify the specific habits and self-regulatory moves students say they use to keep AI from doing the thinking for them. The goal was to surface the practical strategies that turn AI from an answer machine into a partner that still requires evaluation, reflection, and revision. The responses show that a sizable minority already use explicit metacognitive strategies. About a quarter of all respondents talk about verification practices, such as cross-checking AI outputs against textbooks, lecture notes, and other sources, or questioning the answer before accepting it. A similarly large share describes a “tutor, not shortcut” stance, where AI is used for hints, feedback, and explanations but not as a final product. Another group, again close to a quarter, emphasize setting clear learning goals and asking specific questions so that the tool supports a defined purpose rather than replacing their own reasoning. Many also describe mixing AI with traditional study resources, including reading more widely, returning to class materials, or comparing alternative explanations, e.t.c., to force themselves to synthesize. A smaller but visible set ask for policies and training on ethical use, citation, and disclosure.

Taken together, these strategies map onto classic elements of critical thinking: source evaluation, problem framing, and reflective monitoring of one’s own understanding. They also appear to be student-initiated rather than imposed, which suggests that even modest instruction in “critical AI literacy” could amplify practices that are already present. It is important to note that the counts above come from explicit mentions, so they are conservative; students who practice a strategy without naming it would not be captured.

### **RQ3. Differences between quantitative/STEM and qualitative/humanities students**

This question set out to examine how perceived benefits and risks vary across disciplinary cultures. The goal was twofold: first, to compare the kinds of problems students in different fields try to solve with AI, and second, to see whether the risks they notice most often match the demands of their coursework. To make this comparison, fields reported in the open text were grouped into two broad categories. Programs such as accounting, finance, statistics, ICT, mathematics, economics, procurement, and logistics were treated as quantitative/STEM. Programs such as journalism, marketing, linguistics, law, education, management, public administration, and related fields were treated as qualitative/humanities.

Records that were mixed or unclear were not included in the subgroup contrasts. This produced 47 quantitative/STEM records and 55 qualitative/humanities records.

Clear differences appear in the pattern of risks. Among humanities students, 58.2% mention reduced critical thinking, compared with 42.6% in quantitative/STEM, a gap of about 15.6 percentage points. Humanities students also mention internet dependency more often (65.5% versus 48.9%). By contrast, quantitative/STEM students are more likely to mention misinformation or inaccuracy (53.2% versus 36.4%), and they are somewhat more likely to mention over-reliance (27.7% versus 20.0%). These differences align with disciplinary expectations. Courses that hinge on building arguments from sources and synthesizing perspectives are especially sensitive to the risk of accepting ready-made answers, whereas courses that rely on procedural steps and exact calculations are especially sensitive to factual or methodological errors and to a habit of letting the system do the steps. The distribution of use-cases also reflects these emphases. Both groups report using AI for research and information gathering, but quantitative/STEM students more often describe using it to work through problem steps and to translate technical material into simpler language, while humanities students more often highlight brainstorming and support for writing and composition. In practice, this means that the very tasks for which a student turns to AI tend to mirror the risks they notice most: students who ask for worked steps worry about correctness and dependence; students who ask for arguments and phrasing worry about diminished originality and over-accepting a single voice.

The three questions, taken together, paint a consistent picture. Generative AI is a powerful accelerator of baseline understanding, and students can feel that gain very clearly in their day-to-day work. However, the perception that critical thinking may be eroded is widespread enough to be considered a central feature of AI-supported study rather than a marginal concern. A meaningful subset of students already counters this tendency through verification, purposeful prompting, and deliberate combination of AI with other sources, and these habits are concrete targets for instruction and assessment. Finally, the strongest worries shift by discipline in ways that make intuitive sense given the kinds of thinking each field values most. This suggests that guidance should be tuned to local needs: for example, source-triangulation and argument-quality checks in the humanities, and error-analysis and step-validation routines in STEM.

## 5. Conclusion and Recommendations

### 5.1 Conclusion

This study set out to explore how Tanzanian college students perceive and use generative AI (GAI) in their academic practices, particularly focusing on three questions: (1) how GAI supports knowledge construction and whether this comes at the cost of critical thinking, (2) which metacognitive strategies students deploy to sustain critical engagement, and (3) how perceptions differ between quantitative/STEM and qualitative/humanities students. The findings show that GAI has become a mainstream academic tool for research, problem-solving, brainstorming, summarization, and translation. Students appreciate the speed and clarity it provides, describing AI as a scaffold for baseline understanding. Yet nearly half also explicitly recognize the danger of reduced critical thinking.

To counterbalance this risk, many students adopt self-regulatory practices such as cross-verification, purposeful prompting, and combining AI with traditional resources, reflecting an emergent “tutor, not shortcut” mindset. Importantly, disciplinary differences emerged in that STEM students were more concerned with misinformation and over-reliance on step-by-step outputs, while humanities students more frequently cited diminished critical thinking and internet dependency. Together, the results confirm that students see GAI as both a vital accelerator of learning and a potential underminer of intellectual rigor. Therefore, the central tension is not whether AI is useful (it clearly is, as observed from the findings), but how to integrate it without hollowing out the very skills higher education seeks to cultivate.

These findings resonate with and extend the literature in several important ways. First, the efficiency of tradeoff echoes global research that warns of diminished critical thinking when overreliance on AI substitutes for judgment (Barshay, 2025; Pohl et al., 2024). Like European students, Tanzanian learners express worry that outsourcing analysis to machines risks dulling intellectual rigor. Yet unlike their Western peers, they frame this risk within the context of material and infrastructural realities like spotty internet, high data costs, and unreliable electricity (Gondwe, 2024). For instance, one respondent described relying on AI summaries because frequent power outages made sustained deep reading impractical; another noted that purchasing mobile data often meant choosing between extended library research and basic daily needs. In such cases, efficiency is not just a matter of preference but a survival strategy. This necessity amplifies the temptation to adopt shortcuts, even when students are aware that it might compromise their ability to engage in critical reflection. The result is a more complex picture than what emerges in Western contexts: efficiency is simultaneously recognized as a potential intellectual

liability and embraced as a pragmatic response to systemic constraints.

Second, the documentation of student-led metacognitive strategies makes a significant empirical contribution to an emerging but underdeveloped area of literature. The study's findings on “verification practices,” “cross-checking with textbooks,” and the “tutor, not shortcut” mindset provide concrete, on-the-ground evidence for the theoretical concept of “critical AI literacy” (URI ITS, 2024). This argues against a purely dystopian view of student-AI interaction and aligns with the more optimistic perspective of Fischer et al. (2018) on metacognitive scaffolding. However, this study pushes the literature further by demonstrating that these strategies are not just theoretical ideals or the product of top-down instruction; they are organic, bottom-up responses developed by students navigating the risks of GAI on their own. This is a crucial finding because it reveals a latent capacity for critical engagement that can be harnessed and amplified through formal education. Literature has been slow to catalog these spontaneous practices, often focusing on what students *should* do rather than what they *are already doing*. This study argues that pedagogical interventions must be built upon this foundation of existing student agency rather than imposed from scratch.

Third, the clear disciplinary differences validate and geographically diversify a trend recently identified in Western contexts. The finding that humanities students (58.2%) are more concerned about reduced critical thinking than their STEM counterparts (42.6%) perfectly mirror the observations of Pohl et al. (2024) and The Hechinger Report (2025). This strongly suggests that disciplinary epistemology, as in the way knowledge is constructed and validated in different fields, is a more powerful determinant of GAI perception than broad cultural or national context. STEM fields, concerned with “correctness and procedural logic,” naturally breed concerns about “misinformation” and accurate step-by-step solutions. Humanities fields, which “value synthesis, argumentation, and originality,” naturally heighten anxiety about the “loss of independent thought.” This study's contribution is to demonstrate that this disciplinary lens is robust enough to hold true in a Tanzanian university setting, arguing that future research and policy must be disaggregated by discipline to be effective. A one-size-fits-all approach to AI governance is intellectually unsound, as the risks and utility of GAI are inherently field-specific. Finally, by providing rich evidence from Mwanza and Dodoma, this study addresses the major research gap identified in African higher education scholarship (Mare, 2015; Gondwe & Awami, 2025). Unlike much of the Western-focused literature, these findings account for the material realities of digital access, patchwork connectivity, and resource scarcity. They argue for a re-centering of African voices in global debates about GAI in education, which often assume abundant infrastructure and overlook student ingenuity in resource-constrained contexts.

## 5.2 Recommendations

The findings carry significant implications for theory, pedagogy, and policy in African higher education and beyond. Theoretically, this study validates the applicability of theories like cognitive offloading (Risko & Gilbert, 2016) and algorithmic authority (Gillespie, 2014) in a non-Western context but demands their expansion. What we observe in our findings is that the Tanzanian case challenges simplistic accounts of cognitive offloading. Rather than assuming students blindly surrender to algorithmic authority, the evidence points to a dialectic of dependence and resistance. This nuance advances theory by showing that automation bias is not absolute; cultural, pedagogical, and resource contexts mediate how students engage with AI outputs. Therefore, the theories must now account for the "patching" phenomenon, where offloading is a response to infrastructural deficit rather than just a cognitive shortcut. The student response also highlights the role of epistemic agency (the ability to actively manage one's own learning and knowledge validation) as a critical theoretical counterweight to passive acceptance of algorithmic authority.

However, the primary implication is that the goal of education must shift from prohibiting GAI to strategically integrating it while encouraging critical AI literacy. Curricula need explicit modules on prompt engineering, source verification, and ethical citation of AI-generated content. In other words, faculty should tailor AI guidance to disciplinary contexts. Humanities courses might emphasize argument quality, originality checks, and source triangulation, while STEM courses might focus on step validation, error detection, and interpretive reasoning. Assessment must be redesigned to value process over product, emphasizing reflective essays, viva voces, and in-class exercises that demonstrate original thought, thereby making mere submission of AI-generated content insufficient. Furthermore, support must be discipline-specific. For example, humanities courses might teach students to use GAI to generate counterarguments to critique, while STEM courses might focus on using it to explain errors in faulty code or problem-solving approaches.

Most of all, universities, especially those situated in Africa, must move quickly to develop clear, pragmatic, and educational (rather than purely punitive) AI policies. These policies should acknowledge the legitimate academic use of GAI while defining academic misconduct in the age of AI. They must also address the digital divide by ensuring that all students have equitable access to these tools, preventing a new form of inequality between those who can afford premium AI and those who cannot. Finally, further studies should deepen exploration of metacognitive strategies in African contexts, with comparative designs across multiple

regions. Longitudinal research could examine whether "tutor, not shortcut" mindsets sustain over time or fade with increasing familiarity. Future research should, therefore, attempt to conduct follow-up studies to track whether the perceived risks of eroded critical thinking materialize into measurable skills deficits over time; design and test the effectiveness of specific pedagogical interventions (e.g., a specific verification protocol) on mitigating over-reliance and improving learning outcomes, and expand this research to a larger, multi-institutional sample across Sub-Saharan Africa to identify broader regional patterns and nuances.

## References

- Barshay, J. (2025). Proof points: AI in the classroom can hinder deep learning. *The Hechinger Report*.
- Chitanana, T. (2024). Digital literacy in the age of surveillance: A Southern African perspective. *Journal of African Media Studies*, 16(1), 45-62.
- Fischer, C., Pardos, Z. A., & Baker, R. S. (2018). Mining big data in education: Affordances and challenges. *Review of Research in Education*, 42(1), 130-160.
- Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society* (pp. 167–194). MIT Press.
- Gondwe, G. (2024). Digital natives, digital activists in non-digital environments: How the youth in Zambia use mundane technology to circumvent government surveillance and censorship. *Technology in Society*, 79, 102741.
- Gondwe, G. (2025). Perceptions of AI-driven news among contemporary audiences: A study of trust, engagement, and impact. *AI & Society*. Advance online publication.
- Gondwe, G. (2025). Artificial intelligence, journalism, and the Ubuntu robot in sub-Saharan Africa: Towards a normative framework. *Digital Journalism*, 13(4), 826–844.
- Gondwe, G., & Awami, S. (2025). Bridging the knowledge divide: Journalism education and industry needs in East and Southern Africa. *Journalism & Mass Communication Educator*, 80(2), 135–155.
- Kasneji, E., Sessler, K., Kuchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Gunnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., & Schmidt, A.

- (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Mambile, C., & Mwogosi, A. (2025). Transforming higher education in Tanzania: Unleashing the true potential of AI as a transformative learning tool. *Technological Sustainability*, 4(1), 51–76. <https://doi.org/10.1108/TECHS-03-2024-0014>
- Mare, A. (2015). Facebook, youth and political action: A comparative study of Zimbabwe and South Africa. *Unpublished PhD thesis*. Rhodes University, Grahamstown.
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, and abuse. *Human Factors*, 39(2), 230–253. <https://doi.org/10.1518/001872097778543886>
- Pohl, H., Wambsganss, T., & Koedinger, K. R. (2024). *The double-edged sword of generative AI in higher education: A European student perspective*. Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems.
- Rahman, M. M., & Watanobe, Y. (2023). ChatGPT for education and research: A review of benefits and risks. *Sustainability*, 15(12), 9564. <https://doi.org/10.3390/su15129564>
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Ruggiero, T. E. (2000). Uses and gratifications theory in the 21st century. *Mass Communication and Society*, 3(1), 3–37. [https://doi.org/10.1207/S15327825MCS0301\\_02](https://doi.org/10.1207/S15327825MCS0301_02)
- Sánchez-Ruiz, L. M., Moll-López, S., Nuñez-Pérez, A., & Morano-Fernández, J. A. (2023). Generative AI and engineering education: Challenges and opportunities in a new digital era. *IEEE Transactions on Education*, 66(5), 450–461. <https://doi.org/10.1109/TE.2023.3272496>
- Sarakikya, E., & Kitula, L. (2024). The fourth industrial revolution and its impact on Tanzanian higher education. *Tanzania Journal of Development Studies*, 22(1), 34–50.
- Sundar, S. S., & Limperos, A. M. (2013). Uses and gratifications 2.0: New gratifications for new media. *Journal of Broadcasting & Electronic Media*, 57(4), 504–525. <https://doi.org/10.1080/08838151.2013.845827>
- University of Rhode Island, Information Technology Services. (2024). *AI tools: A student introduction to academic uses of AI* [Web resource]. <https://events.uri.edu/event/ai-tools>
- Wineburg, S., & McGrew, S. (2017). *Lateral reading and the nature of expertise: Reading less and learning more when evaluating digital information* (Working Paper). Stanford History Education Group. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3048994](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3048994)
- 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, 39. <https://doi.org/10.1186/s41239-019-0171-0>