iYunivesithi Walter Sisulu | Innovative Journal of African Education

Generative AI and Cognitive Autonomy in Cameroonian Higher Education: A Conditional Amplification Framework for Learning

Rene Blaise Mbop Ackok
Adventist University Cosendai
Higher Institute of Educational Sciences
Nanga-Eboko, Cameroon
Email: sakanemayi@yahoo.fr
***ORCID iD:** 0009-0003-7171-9485*

Abstract

Generative artificial intelligence (GAI) is rapidly entering higher education, yet its implications for learner autonomy remain unsettled, particularly in resource-constrained contexts. This explanatory sequential mixed-methods study examined how pedagogical GAI use relates to self-regulated learning, cognitive engagement, augmented cognitive autonomy (ACA), and deep learning in six Cameroonian higher education institutions. Quantitative data from 812 students and faculty members were analysed using partial least squares structural equation modelling, followed by 24 semi-structured interviews. GAI use was positively associated with self-regulation, while self-regulation and cognitive engagement sequentially mediated its relationship with ACA. ACA was strongly associated with deep learning. AI literacy and institutional governance strengthened the relationship between GAI use and ACA. Qualitative findings revealed both reflective, dialogic uses and cognitive delegation. The findings support an emerging Augmented Cognitive Autonomy framework in which GAI amplifies learning conditionally, depending on users' regulatory capacities and institutional support, with implications for equity and human capability development.

Keywords: generative artificial intelligence; augmented cognitive autonomy; self-regulated learning; AI literacy; Cameroon

1. Introduction and Background

Generative artificial intelligence (GAI) has moved quickly from a specialist technology to an everyday academic resource. Students and academics now use systems such as ChatGPT, Gemini, Claude and Copilot to search for information, explain difficult concepts, generate examples, draft text, test ideas and obtain feedback. Unlike earlier educational technologies that primarily expanded access to information or communication, generative systems participate directly in activities that have traditionally signalled intellectual work: synthesis, argumentation, problem solving and knowledge production (Dwivedi et al., 2023; Kasneci et al., 2023; Tlili et al., 2023). This shift creates a difficult educational question. If an artificial agent can perform part of the reasoning involved in a task, what does it mean for a learner to remain intellectually autonomous?

©Authors
This work is licensed under a Creative Commons Attribution-NonCommercial-No Derivatives 4.0 International License (CC BY-NC-ND 4.0).

Research has so far produced two broad narratives. One emphasises the potential of GAI to personalise support, provide immediate feedback, widen access to explanations and scaffold learning (Chan & Hu, 2023; UNESCO, 2023). The other highlights cognitive offloading, reduced effort, unreliable outputs, academic integrity concerns and the possibility that learners may substitute generated answers for understanding (Lodge et al., 2023; Perkins et al., 2024). Treating these narratives as mutually exclusive is increasingly unhelpful. The same technology can prompt a learner to compare explanations and interrogate evidence, or allow that learner to bypass the very reasoning a task was designed to develop. The educational effect therefore depends not only on access to GAI but on how the interaction is regulated.

This issue is especially important for learning and development. Digital technologies can expand educational opportunity while reproducing or deepening inequalities when access, skills and institutional support are uneven (Helsper, 2021; van Dijk, 2020). In many African higher education settings, GAI is being adopted within institutions that differ substantially in connectivity, digital support, staff preparation and formal guidance. Cameroon offers a useful setting in which to examine these dynamics because public, private and faith-based institutions coexist within a heterogeneous sociotechnical environment. The point is not to treat Africa as a single or deficient context, but to examine how a rapidly diffusing technology interacts with unequal educational resources and institutional capacities.

Cameroon illustrates these dynamics with particular clarity. The country's bilingual (French–English) higher education landscape, its coexistence of state-funded, private and confessional universities, and its uneven distribution of broadband access and institutional digital-support services together produce a natural variation in the conditions under which students and faculty encounter generative systems. Some institutions have issued explicit guidance on academic integrity and AI use; others have not yet formulated a clear position. Some students access GAI chiefly through mobile data on personal handsets, with the attendant costs and connectivity constraints this entails, while others rely on institutional computer laboratories or personally owned laptops with more stable connections. This heterogeneity is not incidental to the research question; it is precisely the kind of structural variation needed to test whether AI literacy and institutional governance condition the relationship between GAI use and cognitive autonomy, rather than assuming, from contexts of near-universal high-speed access, that this relationship is uniform. Examining Cameroon therefore serves both a contextual and an analytical purpose: it documents an under-studied setting and supplies the variation the conceptual model requires.

Three gaps motivate the study. First, much of the early GAI literature documents use, perceptions, opportunities and risks rather than the mechanisms through which GAI may alter learning processes. Second, autonomy is often discussed as either preserved or lost, without sufficient attention to the regulatory work required to think with an artificial agent while retaining judgement and responsibility. Third, evidence from African higher education remains comparatively limited even though these settings can illuminate how AI literacy and institutional conditions shape technology-enabled learning.

We therefore examine GAI not as an inherently emancipatory or harmful technology but as a potential cognitive amplifier whose educational consequences are conditional. We use the term augmented cognitive autonomy (ACA) to describe a learner's capacity to

use an artificial agent strategically while retaining control of reasoning, verification, metacognitive monitoring and intellectual responsibility. The study asks: Under what conditions is pedagogical GAI use associated with stronger cognitive autonomy and deep learning in Cameroonian higher education, and how do self-regulation, cognitive engagement, AI literacy and institutional governance help explain this relationship?

The study makes three contributions. Empirically, it combines structural modelling with qualitative accounts from a context underrepresented in international GAI research. Conceptually, it develops and tests ACA as a multidimensional construct rather than assuming that technology use and autonomy are opposites. For learning for development, it shifts attention from access alone to the capacity to orchestrate human-AI interaction critically, a capacity that may itself become a new dimension of educational inequality.

2. Literature Review

2.1 Generative AI, self-regulation and learner autonomy

The emerging empirical literature no longer supports a simple opposition between generative AI and learner autonomy. Generative systems can provide immediate explanation, feedback, practice and adaptive support, but these affordances become educationally meaningful only when learners continue to plan, monitor and evaluate their own activity. Chiu (2024) accordingly conceptualises ChatGPT use through self-determination and self-regulated learning, emphasising that the same technological affordance can support or weaken autonomous regulation depending on how it is incorporated into learning. Wu and Chiu (2025) similarly show that self-regulated learning with generative AI is configurational: productive outcomes arise from combinations of learner characteristics and AI affordances rather than from tool use alone. This literature suggests that frequency of use is an insufficient explanatory variable; the regulatory quality of human–AI interaction is the more consequential object of study.

2.2 Cognitive offloading, verification and productive struggle

A parallel literature warns that generative AI can relocate cognitive work from the learner to the system. Cognitive offloading is not inherently detrimental: external resources have always supported memory, calculation and problem solving. The educational risk arises when delegation replaces the effort through which understanding, monitoring and evaluative judgement are developed. Tate et al. (2025), in a writing-intensive undergraduate course, demonstrate that generative AI can be integrated without off-loading learning when use is deliberately scaffolded around understanding, corroboration and responsible incorporation of generated material. In Cameroon, Moluayonge (2025) reports an empirical relationship between AI use, cognitive offloading and critical-thinking development, underscoring the relevance of this problem within the national higher-education context. These findings make verification, source comparison and intellectual ownership central to any account of AI-mediated autonomy.

Recent evidence sharpens this conditional account. A 2026 systematic review of 67 empirical higher-education studies found that ChatGPT was more consistently associated with gains in metacognitive regulation, argumentative reasoning and higher-order thinking when use was scaffolded and inquiry-oriented, whereas unstructured use was more often associated with cognitive offloading and diminished critical engagement (Li et al., 2026). Longitudinal prompt analysis likewise shows that student–AI collaboration

can drift toward low-effort, minimally contextualised prompting and acceptance of model output, while higher-quality work is associated with more deliberate contextualisation and task guidance (Misiejuk et al., 2026). These findings strengthen the need to distinguish access to GAI from the learner's capacity to regulate what is delegated, verified and retained as human intellectual responsibility.

2.3 AI literacy, institutional conditions and African higher education

AI literacy is increasingly understood as more than technical familiarity with a tool. It includes the capacity to recognise limitations, evaluate outputs, make informed choices about delegation and use AI responsibly. Yet these capacities develop within institutional environments. Rules on disclosure, assessment design, staff preparation, access, connectivity and formal guidance can either support or constrain critical use. Evidence from African higher education remains thinner than evidence from North America, Europe and parts of Asia, while preliminary multi-country work in African health-professions education reports widespread use alongside limited formal training and institutional guidance. The African question is therefore not simply whether students have access to generative AI, but whether institutions create the pedagogical and governance conditions under which access can become educational capability.

2.4 Gap addressed by the present study

Three gaps remain. First, many studies examine attitudes, adoption, perceived usefulness or broad learning outcomes without modelling the regulatory mechanisms linking generative AI use to autonomy. Second, self-regulation, critical thinking, learner autonomy and AI literacy are often invoked separately, leaving unclear how responsibility is maintained when cognition is distributed between a human learner and a generative system. Third, relatively little empirical work tests these mechanisms in heterogeneous African higher-education environments. The present study addresses these gaps by examining a conditional pathway from pedagogical GAI use through self-regulation and cognitive engagement to an emerging construct—Augmented Cognitive Autonomy (ACA)—while testing AI literacy and institutional governance as contextual moderators and using qualitative evidence to distinguish cognitive augmentation from cognitive delegation.

3. Theoretical Framework and Hypotheses

Self-regulated learning provides a first foundation for understanding productive GAI use. Self-regulation involves setting goals, selecting strategies, monitoring progress and adapting action in response to feedback (Panadero, 2017; Zimmerman, 2002). GAI can support these processes when learners use it to clarify instructions, generate practice questions, diagnose gaps or request alternative explanations. Yet the same system can weaken regulation if it assumes the learner's role in defining the problem, selecting a strategy, producing the reasoning and judging the answer. Frequency of use is therefore less informative than the regulatory quality of use. This reframing carries a methodological implication as much as a theoretical one: studies that measure only frequency of use or general attitude toward GAI are likely to under-specify the mechanism through which the technology shapes learning outcomes, which is why the present study models self-regulation and cognitive engagement as sequential mediators rather than as simple covariates.

A second foundation comes from distributed cognition. Cognition has long been distributed across people, artefacts and symbolic resources rather than residing solely within an individual mind (Hutchins, 1995; Salomon, 1993). Books, calculators, search engines and statistical software all extend human cognitive activity. GAI intensifies this relationship because it does not merely store or retrieve information; it generates plausible explanations, arguments and solutions. This makes it useful to distinguish a cognitive partner from a cognitive substitute. A partner extends the learner's activity while the learner retains goals, evaluative criteria and responsibility. A substitute performs those functions in place of the learner. This partner–substitute distinction echoes, without collapsing into, Vygotsky's (1978) account of mediated action, in which tools and signs reorganise rather than simply replace human mental functioning. On this view, the pedagogical question is not whether a tool mediates cognition, since all higher mental functions are in some sense mediated, but whether the learner remains the author of the goals and standards that guide the mediated activity.

Agency and self-determination add a third dimension. Agency concerns intentional action and perceived control, while autonomy involves acting with volition and ownership (Bandura, 1986; Deci & Ryan, 2000). In AI-mediated learning, autonomy cannot reasonably mean independence from all technological mediation. A more useful question is whether the learner remains able to direct, interrogate and take responsibility for a distributed cognitive process. This is consistent with self-determination theory's distinction between autonomous and controlled forms of regulation: a learner who selects and steers AI assistance in service of self-endorsed goals is exercising autonomous regulation, whereas a learner who defers to AI output chiefly because it is easier or expected is exhibiting a more externally controlled form of engagement, even where the surface behaviour, prompting a chatbot, looks similar in both cases.

We integrate these perspectives in an emerging Augmented Cognitive Autonomy (ACA) framework. ACA comprises five related capacities: cognitive control, critical verification, metacognitive regulation, intellectual responsibility and strategic use of GAI. The framework proposes conditional cognitive amplification: GAI can extend learning when these capacities are activated, but can also facilitate cognitive delegation when they are weak. AI literacy and institutional governance are treated as contextual resources because users need both an understanding of AI's capabilities and limitations and environments that provide expectations, guidance and support (Long & Magerko, 2020; UNESCO, 2023).

A construct such as ACA invites an immediate and reasonable question: does it add explanatory value beyond established constructs such as self-regulation, critical thinking or AI literacy, or does it simply relabel existing variables under a new name? We treat this as an empirical rather than a rhetorical matter. Conceptually, ACA is distinguished from general self-regulation by its explicit focus on a human–AI dyad rather than on learning behaviour in the abstract: its items assess whether a person can direct, verify and take responsibility for output that another agent produced, a demand that did not exist in this specific form before generative systems could produce fluent, plausible text at scale. It is distinguished from AI literacy by locating the construct in situated cognitive practice, that is, in what a learner actually does within a given interaction, rather than in general knowledge about how AI systems are built or trained; a learner can, in principle, understand how a large language model is trained while still delegating uncritically under time pressure. The measurement results reported below, in which

ACA shows adequate discriminant validity from the AI-literacy and governance indicators via heterotrait-monotrait ratios, offer preliminary rather than conclusive evidence for this distinction, and we return to the point in the Limitations section.

ACA is therefore not proposed as a synonym for self-regulation, critical thinking, learner autonomy or AI literacy. Self-regulation concerns the management of one's learning process; critical thinking concerns the evaluation of claims and reasons; learner autonomy concerns volitional ownership of learning; and AI literacy concerns knowledge and competence for understanding and using AI. ACA targets a narrower interactional problem: whether a person can strategically distribute cognitive work to a generative agent while retaining control over goals, verification, metacognitive monitoring, final judgement and intellectual responsibility. Its proposed incremental value is thus relational and situated. This distinction remains provisional and requires independent tests of incremental and discriminant validity.

Table 1

Conceptual Boundaries Between ACA and Adjacent Constructs

Construct	Primary focus	Relation to AI	What it explains	Boundary relative to ACA
Self-regulated learning	Planning, monitoring and adapting learning	AI may support or disrupt regulation	How learners manage goals, strategies and progress	ACA concerns regulation of delegated cognitive work within a human–AI interaction.
Critical thinking	Evaluation of claims, evidence and reasoning	AI outputs become objects of scrutiny	Quality of judgement and evaluation	Critical verification is one ACA capacity, but ACA also includes strategic delegation, metacognition and responsibility.
Learner autonomy	Volition, ownership and self-directed action	AI can support or constrain autonomous action	Who owns and directs the learning process	ACA operationalises autonomy specifically when cognitive activity is shared with a generative agent.
AI literacy	Knowledge, skills and dispositions for understanding and using AI	AI is the object of literacy	Whether users understand capabilities, limits and responsible use	AI literacy is treated as a contextual resource; knowledgeable users may still delegate uncritically in a specific task.

Construct	Primary focus	Relation to AI	What it explains	Boundary relative to ACA
Augmented Cognitive Autonomy (ACA)	Responsible orchestration of human–AI cognition	AI is a cognitive partner whose contribution must be directed and verified	Whether users retain control, verification, metacognitive monitoring, judgement and responsibility while using GAI	Emerging situated construct requiring independent validation and incremental-validity testing.

This formulation also has an equity implication. Digital inequality is not limited to devices and connectivity; it includes differences in skills and in the benefits people can derive from technology (Helsper, 2021; van Dijk, 2020). With GAI, a further inequality may emerge between users who can critically orchestrate human-AI interaction and those who mainly consume generated outputs. We refer to this as cognitive orchestration inequality. The term is offered as an interpretive proposition requiring further validation, not as an established universal category.

Four hypotheses were tested:

H1. Pedagogical use of GAI is positively associated with self-regulated learning.

H2. Self-regulated learning and cognitive engagement sequentially mediate the relationship between pedagogical GAI use and ACA.

H3. ACA is positively associated with deep learning.

H4. AI literacy and institutional governance positively moderate the relationship between pedagogical GAI use and ACA.

3. Research Design and Methods

Research Design

An explanatory sequential mixed-methods design was used within a pragmatic orientation (Creswell & Plano Clark, 2018; Morgan, 2014). The quantitative phase tested the proposed relationships among GAI use, self-regulation, cognitive engagement, ACA and deep learning. The subsequent qualitative phase explored how participants experienced cognitive support, critical engagement and delegation when using GAI. Integration occurred at interpretation through meta-inferences that compared the structural results with recurring qualitative patterns. Partial least squares structural equation modelling (PLS-SEM) was selected over covariance-based SEM because the primary analytic objective was prediction-oriented, explaining variance in ACA and deep learning, rather than confirmatory testing of a single, fully specified covariance structure, and because the model combines a newly developed construct (ACA) with established reflective measures drawn from prior scales, a combination for which PLS-SEM is well suited (Hair et al., 2022).

Context and Participants

The study was conducted in six Cameroonian higher education institutions representing public, private and faith-based sectors. A stratified purposive sampling strategy was used to obtain variation across institutional sectors and participant profiles, so that no single sector or discipline would dominate the pool through convenience effects. The quantitative sample comprised 812 participants: 684 students (84.2%) and 128 faculty members (15.8%). Women represented 47.3% of the sample and men 52.7%; mean age was 28.7 years (SD = 8.6). Participants represented humanities and social sciences (34.6%), education (26.6%), business and economics (17.6%), and science and technology (21.2%). Because recruitment relied on institutional gatekeepers and voluntary response rather than a documented probability-sampling procedure, the sample is not treated as nationally representative and population-level generalisations are not claimed; it is treated as analytically rather than statistically representative of GAI use in Cameroonian higher education.

The qualitative phase purposively selected 24 participants (16 students and eight faculty members) for explanatory follow-up. Selection sought contrast in participant role and reported patterns of GAI use so that the interviews could illuminate why similar access was associated with different forms of learning behaviour. The qualitative phase was not designed to reproduce the demographic distribution of the survey or to support population estimates. Age- and gender-disaggregated characteristics of the interview subsample were not retained in the manuscript dataset available for this revision and are therefore not inferred or reconstructed.

Measures

The quantitative questionnaire used seven-point Likert response scales (1 = strongly disagree to 7 = strongly agree). Measures of pedagogical GAI use, self-regulated learning, cognitive engagement, deep learning, AI literacy and institutional governance were adapted from established literature, including Zimmerman (2002), Panadero (2017), Fredricks et al. (2004), Biggs and Tang (2011), Long and Magerko (2020), Chan and Hu (2023), and UNESCO (2023).

ACA was developed for this study following established scale-development principles (Boateng et al., 2018; DeVellis & Thorpe, 2021). An initial pool of 28 items was derived from literature on distributed cognition, self-regulation, critical thinking and educational AI. Six experts reviewed content, reducing the pool to 20 items. A pretest with 60 participants informed wording revisions. Exploratory factor analysis suggested five dimensions accounting for 68.4% of variance. The final ten-item solution represented cognitive control, critical thinking, metacognitive regulation, intellectual responsibility and strategic GAI use. Standardised loadings ranged from .79 to .88; AVE was .61, Cronbach's alpha .89 and composite reliability .92 during scale development.

Data Collection and Analysis

Quantitative data were collected from January to March 2026 through online and in-person administration. Semi-structured interviews followed from April to May 2026. Interviews were recorded with participants' consent and transcribed. The quotations reported in the findings are authentic excerpts from the interview corpus, translated into English for publication while preserving their original meaning. Participation was voluntary and based on informed consent; confidentiality and anonymisation were

maintained in reporting. Ethical approval was obtained from the Ethics Committee of the Higher Institute of Educational Sciences, Adventist University Cosendai, Cameroon.

Quantitative analyses were conducted in SmartPLS 4. The measurement model was assessed using internal consistency, composite reliability, average variance extracted, heterotrait-monotrait ratios and variance inflation factors. The structural model was examined through path coefficients, mediation and moderation effects, effect sizes, R^2 and Q^2, consistent with current PLS-SEM reporting recommendations (Hair et al., 2022). Statistical significance was estimated using 5,000 bootstrap resamples with bias-corrected confidence intervals. SRMR and NFI are reported as supplementary approximate-fit indicators rather than treated as decisive evidence of model validity.

Interview transcripts were analysed using reflexive thematic analysis (Braun & Clarke, 2022). Analysis proceeded through repeated familiarisation with the transcripts, initial coding, development and refinement of candidate patterns, and interpretive review against the explanatory purpose of the qualitative phase. Coding attended particularly to planning, verification, dialogue with GAI, changes in effort, intellectual responsibility and institutional support. Themes were treated as analytically constructed patterns rather than as mechanically counted categories. Quantitative and qualitative findings were then integrated through a joint display to identify convergence, complementarity and tension. Because the study is cross-sectional, all structural relationships are interpreted as associations consistent with the proposed framework rather than proof of causality.

Table 2

Characteristics of the Quantitative Sample (N = 812)

Characteristic	Category	n	%
Gender	Women	384	47.3
	Men	428	52.7
Age	18–24 years	302	37.2
	25–34 years	318	39.2
	≥35 years	192	23.6
Status	Students	684	84.2
	Faculty	128	15.8
Institution	Public	356	43.8
	Private	278	34.2
	Faith-based	178	21.9
GAI experience	Low	215	26.5
	Moderate	384	47.3
	High	213	26.2

4. Findings

Participant Use of GAI and Measurement Model

GAI use was widespread: 82.5% of respondents reported using generative AI at least several times per week and 27.1% reported daily use. The most common purposes were information seeking (84.7%), explaining difficult concepts (76.3%) and academic writing support (69.5%). The measurement model showed satisfactory reliability and convergent validity. Cronbach's alpha ranged from .87 to .93, composite reliability from .91 to .94, and AVE from .61 to .69. HTMT values ranged from .52 to .82, below the conservative .85 criterion. VIF values ranged from 1.24 to 2.63, indicating no problematic collinearity in the reported model. Q^2 values were positive for self-regulation (.22), cognitive engagement (.31), ACA (.38) and deep learning (.33).

Structural Model

The model accounted for 34% of the variance in self-regulation, 45% in cognitive engagement, 58% in ACA and 47% in deep learning. The reported SRMR was .048 and NFI .917.

H1 was supported: pedagogical GAI use was positively associated with self-regulated learning ($\beta = .58$, $t = 17.11$, $p < .001$, 95% CI [.51, .64], $f^2 = .50$). Self-regulation was positively associated with cognitive engagement ($\beta = .67$, $p < .001$), and cognitive engagement with ACA ($\beta = .63$, $p < .001$). The sequential indirect effect from GAI use through self-regulation and cognitive engagement to ACA was significant ($\beta = .25$, $t = 7.84$, $p < .001$, 95% CI [.18, .33]), supporting H2.

ACA was strongly associated with deep learning ($\beta = .68$, $t = 22.15$, $p < .001$, 95% CI [.62, .74], $f^2 = .71$), supporting H3. AI literacy strengthened the association between GAI use and ACA ($\beta = .17$, $t = 4.62$, $p < .001$, 95% CI [.09, .25], $f^2 = .06$), while institutional governance showed a smaller positive moderation effect ($\beta = .12$, $t = 2.98$, $p = .003$, 95% CI [.04, .20], $f^2 = .04$). H4 was therefore supported, although both moderation effects were modest.

Qualitative Findings

Five recurring patterns helped explain the structural results. First, participants described GAI as metacognitive scaffolding. Students used repeated explanations, self-testing and reformulation to identify gaps and adjust study strategies. One student explained, “I use it mainly to understand what I did not understand in class. I ask it to rephrase until it becomes clear.” This pattern is consistent with the association between GAI use and self-regulation.

Second, experienced users described a dialogic form of distributed cognition. Rather than accepting an answer, they challenged the system, requested alternatives and compared outputs with other sources. A participant noted, “I never take an answer as it is. I dialogue with the AI, challenge it, ask for references and compare with other sources.” In these accounts, GAI extended reasoning while the user retained evaluative control.

Third, other participants described cognitive delegation. They reported searching less independently, asking for answers before attempting problems, or producing text more quickly without being confident that understanding had improved. These accounts qualify any simple interpretation of high GAI use as educational benefit.

Fourth, participants repeatedly described unequal capacity to benefit from GAI. Faculty observed that students who knew how to formulate questions, verify claims and use

outputs strategically appeared to learn more than those seeking rapid answers. One student captured the institutional dimension: “We all use AI, but nobody has taught us how to use it properly.” These accounts informed the proposition of cognitive orchestration inequality.

Finally, several interviews were explicitly ambivalent. Participants valued time savings and personalised explanation while also recognising the temptation to reduce effort. The same user could move between augmentation and delegation depending on the task, time pressure and verification practices.

Mixed-Methods Integration

The two strands converged on a conditional rather than deterministic account of GAI. The quantitative model linked GAI use to ACA indirectly through self-regulation and cognitive engagement; interviews showed what these mechanisms looked like in practice, including iterative questioning, verification and strategy adjustment. The moderation results indicated that AI literacy and governance mattered; qualitative accounts suggested why, by showing that users often lacked explicit preparation for critical AI use. At the same time, interview evidence of cognitive delegation cautions against interpreting positive average path coefficients as universal individual trajectories.

Table 3

Measurement and Structural Model Summary

Indicator / Relationship	Estimate	Interpretation
Cronbach's α	.87–.93	Satisfactory internal consistency
Composite reliability	.91–.94	Satisfactory reliability
AVE	.61–.69	Convergent validity supported
HTMT	.52–.82	Discriminant validity supported
GAI use → Self-regulation	$\beta = .58$, $p < .001$	H1 supported
Sequential indirect effect to ACA	$\beta = .25$, $p < .001$	H2 supported
ACA → Deep learning	$\beta = .68$, $p < .001$	H3 supported
AI literacy × GAI use → ACA	$\beta = .17$, $p < .001$	H4 supported; modest effect
Governance × GAI use → ACA	$\beta = .12$, $p = .003$	H4 supported; small effect
R^2 ACA / Deep learning	.58 / .47	Moderate-to-substantial explained variance

Table 4
Mixed-Methods Joint Display

Quantitative finding	Qualitative explanation	Integrated inference
GAI use → self-regulation (β=.58)	Reformulation, self-testing and gap identification	GAI can scaffold regulation when learners actively direct the interaction.
Self-regulation → engagement → ACA	Dialogue, challenge and source comparison	Autonomy is associated with active cognitive engagement, not mere tool use.
AI literacy moderation (β=.17)	Skilled users verify and compare outputs	AI literacy may help convert access into productive cognitive partnership.
Governance moderation (β=.12)	Participants reported little formal preparation	Institutional guidance appears supportive but cannot substitute for learner regulation.
Positive average pathways	Accounts of reduced effort and answer-seeking	GAI effects are heterogeneous; cognitive delegation remains a competing trajectory.

5. Discussion of Findings

The findings suggest that the educational significance of GAI lies less in whether it is used than in how cognitive work is distributed between the learner and the system. This helps move the debate beyond a binary in which GAI either enhances learning or undermines it. In this study, pedagogical GAI use was associated with stronger self-regulation, but the qualitative findings showed that this relationship was not automatic. Some learners used the technology to expose gaps in understanding and test themselves; others used it to shorten or bypass effort. GAI is therefore better understood here as a conditional cognitive amplifier.

This result is consistent with, but more specific than, the emerging international evidence. Recent synthesis indicates that the cognitive consequences of ChatGPT vary with pedagogical scaffolding and learner regulation rather than with tool exposure alone (Li et al., 2026). Likewise, longitudinal analysis of student prompting shows that the distribution of cognitive effort within human–AI interaction changes across tasks and is visible in the quality of prompting and critical engagement (Misiejuk et al., 2026). The present study extends this literature by modelling a named orchestration construct—ACA—and by testing a conditional pathway in which self-regulation and cognitive engagement are associated with ACA while AI literacy and institutional governance moderate the GAI–ACA relationship. The contribution is therefore not the claim that regulation matters, which prior studies increasingly support, but an empirically specified account of how regulation, engagement, literacy and governance jointly relate to retained cognitive control in an African higher-education setting.

This interpretation extends self-regulated learning research. Traditional accounts emphasise learners' planning, monitoring and strategy adjustment (Panadero, 2017; Zimmerman, 2002). GAI adds an interactive agent capable of suggesting goals, strategies and evaluations. The learner's regulatory task consequently changes: autonomy involves not only regulating one's own cognition but also regulating what is delegated to the artificial system. The positive sequential pathway through self-regulation and cognitive

engagement is consistent with this interpretation. However, because the data are cross-sectional, the pathway should be read as evidence compatible with the framework, not as a demonstrated causal sequence.

The findings also refine distributed cognition. External cognitive resources do not necessarily diminish autonomy (Hutchins, 1995; Salomon, 1993). What matters is who controls the goals, standards of evidence and final judgement. The qualitative contrast between dialogue and delegation makes this distinction visible. Participants who questioned, compared and verified AI outputs treated GAI as a cognitive partner. Participants who requested completed answers without prior effort or verification described a form of substitution. ACA captures this boundary by defining autonomy as responsible orchestration rather than technological independence.

For this reason, the evidence supports an emerging Augmented Cognitive Autonomy framework, but it would be premature to claim a fully established theory. The present study offers an empirically informed model and a newly developed measure that require replication, longitudinal testing and cross-cultural validation. This more cautious formulation strengthens rather than weakens the contribution: it specifies a testable proposition about autonomy in AI-mediated learning and makes clear what future evidence would be needed.

The moderation findings are particularly relevant to learning for development. AI literacy had a statistically significant but modest strengthening effect, and institutional governance had a smaller one. These results indicate that access to GAI alone is unlikely to equalise educational opportunity. Learners also need capabilities to evaluate outputs, recognise limitations, formulate productive prompts, compare sources and decide when not to delegate. Institutions need guidance that goes beyond policing academic misconduct to include pedagogical support for critical use. This aligns with broader digital inequality research, which shows that differences in skills and outcomes persist even after basic access improves (Helsper, 2021; van Dijk, 2020).

The concept of cognitive orchestration inequality emerges from this intersection. It describes differences in people's capacity to convert access to GAI into intellectually productive interaction. The interviews provide initial evidence for the phenomenon, while the moderation results suggest that AI literacy and governance are relevant resources. The concept should nevertheless be treated as a proposition rather than a confirmed new form of inequality. Future work should test whether orchestration capacity predicts learning outcomes after accounting for socioeconomic background, connectivity, prior achievement and disciplinary context.

The present pattern is consistent with, but more specific than, recent work showing that generative AI can support self-regulated learning when it is embedded in deliberate regulatory practices (Chiu, 2024; Wu & Chiu, 2025). It also converges with instructional research showing that AI-supported writing need not produce cognitive offloading when learners are required to understand, corroborate and actively incorporate generated material (Tate et al., 2025). The contribution of the present model is to connect these two literatures: it proposes that the educational value of GAI depends on whether regulatory activity and cognitive engagement culminate in retained human control over verification and responsibility. The Cameroonian evidence is particularly consequential because local work has already identified cognitive offloading as a mechanism linking AI use and critical-thinking outcomes (Moluayonge, 2025). ACA therefore attempts to specify the

interactional capacities that may separate productive externalisation from uncritical delegation.

Boundary Conditions and Alternative Interpretations

Several alternative readings of these findings deserve explicit consideration. A first possible objection is that framing GAI as an *amplifier* risks a subtly technologically deterministic narrative, in which the technology is assumed to have a direction of effect that merely awaits the right moderating conditions. The present framework was deliberately built to resist this reading: amplification is modelled as conditional rather than fixed, the qualitative material documents dialogue and delegation coexisting within the same institutions and often within the same individuals across tasks, and the structural model treats AI literacy and governance as moderators precisely because the sign and strength of the GAI–ACA relationship is not assumed to be constant across users.

A second objection concerns the risk of psychologising a structural problem, that is, of locating responsibility for cognitive delegation in individual learners' regulatory capacity while underplaying infrastructural and institutional inequality. This is a legitimate concern in a Cameroonian setting where connectivity, device access and formal AI guidance are unevenly distributed. The inclusion of institutional governance as a moderator, and the proposition of cognitive orchestration inequality as a structural rather than a deficit category, are intended precisely to keep responsibility distributed across learners and institutions rather than resting it solely on individual self-discipline. We nonetheless acknowledge that the ACA measure itself is administered at the individual level, and that a purely individual-level instrument cannot, on its own, settle how much of the observed variance is attributable to structural conditions rather than personal disposition; this is a boundary condition of the design rather than of the construct, and it is a natural target for multilevel extensions of the model.

A third objection concerns the cultural specificity of autonomy as a construct. ACA draws on theoretical traditions, self-determination theory and social-cognitive theory in particular, that were developed largely within Western, individualist psychological research (Bandura, 1986; Deci & Ryan, 2000). Higher education in Cameroon, as in many African settings, also involves collaborative and communal epistemic practices, including peer study groups and shared verification of information, that an individually administered scale may not fully capture. We treat this as an open question rather than a settled matter: the present instrument measures an individual's capacity to direct and verify AI-assisted work, but it does not test whether collective or dialogic forms of verification, several students jointly interrogating an AI-generated answer, for instance, function as an equivalent or even superior route to the same underlying capacities. Cross-cultural and qualitative work that examines collective orchestration of GAI use, rather than assuming the individual as the sole unit of analysis, is an important next step.

A fourth objection is that a framework built around conditionality risks becoming unfalsifiable, able to accommodate any pattern of results by invoking unmeasured moderators. We therefore state explicitly what would count as evidence against the framework as currently specified: a null or negative moderation effect of AI literacy on the GAI–ACA relationship, a direct effect of GAI use on ACA that was not substantially reduced once self-regulation and cognitive engagement were entered as mediators, or a

failure of the ACA measure to show discriminant validity from AI literacy in independent samples would each cast doubt on the specific mechanism proposed here, even if a general association between GAI use and learning outcomes persisted. Framing the theory in this falsifiable form is intended to make the present study a starting point for cumulative testing rather than a closed interpretive account.

The Cameroonian setting sharpens these implications. The six participating institutions represented different sectors and participants reported heterogeneous levels of experience with GAI. In such environments, rapid adoption without corresponding institutional preparation may widen differences between users who can critically exploit generative systems and those who mainly consume their outputs. Development-oriented technology policy should therefore ask not only who can access AI, but who can use it to expand human capability. This reframes AI literacy as an educational equity issue.

For teaching practice, the results favour designs that make cognitive responsibility visible. Academics can require students to compare AI-generated explanations with scholarly sources, document how outputs were checked, identify errors or omissions, explain why a prompt was revised, and distinguish tasks where AI support is appropriate from tasks where independent reasoning should come first. Such practices use GAI without allowing the generated product to become a proxy for learning.

For institutions, governance should combine clear expectations with capacity building. Policies focused only on prohibition or detection may miss the central learning problem. Students and staff need shared guidance on verification, disclosure, privacy, intellectual responsibility and appropriate delegation. Faculty development is equally important because teachers must redesign tasks and assessment so that judgement, reasoning and reflection remain observable even when generative tools are available.

Limitations and Future Research

Several limitations constrain interpretation. First, the cross-sectional design cannot establish causal effects, despite the directional structure of the tested model. Longitudinal and experimental studies are needed to examine whether GAI use changes self-regulation and autonomy over time, and whether the sequential mediation pathway identified here holds when measured across repeated occasions rather than inferred from a single wave of data. Second, the main variables were self-reported and may be affected by social desirability, common-method variance and differences between perceived and demonstrated competence; a learner's account of verifying AI output, for example, may not correspond exactly to observable verification behaviour. Future research should incorporate performance tasks, trace data and direct measures of verification behaviour alongside self-report. Accordingly, the statistically significant indirect pathway is described as an indirect association within the specified model, not as temporal evidence that self-regulation and engagement causally transmit an effect of GAI use to ACA.

Third, students and faculty members were included in the same quantitative model because the focal constructs were conceptualised at the individual level as patterns of AI-mediated cognitive activity. However, the two groups occupy different roles in higher education, and the present study did not test measurement invariance or compare structural paths between students and faculty. The pooled estimates should therefore be interpreted as overall associations within the participating sample rather than evidence that the model operates identically across both groups. Future research should test the

framework separately across learner and educator populations and conduct formal multigroup analyses.

Fourth, ACA is a newly developed construct. Its five-dimensional structure and ten-item measure showed encouraging initial psychometric properties, but external validation, test-retest evidence and cross-cultural measurement invariance are still required, including outside francophone-anglophone Cameroon and outside higher education altogether. The concepts of conditional cognitive amplification and cognitive orchestration inequality should likewise be regarded as emerging propositions to be tested rather than settled theoretical facts. A decisive next test is incremental validity: ACA should predict relevant learning outcomes after controlling for established measures of self-regulated learning, critical thinking, learner autonomy and AI literacy. Failure to add explanatory value beyond these adjacent constructs would weaken the case for ACA as a distinct construct.

Fifth, the qualitative subsample of 24 participants, while adequate for the explanatory purpose of this phase, cannot support claims of thematic saturation across the full diversity of institutions, disciplines and GAI tools represented in the wider sample; the qualitative findings are best read as illustrative of mechanisms rather than as a representative account of all participants' experiences. Finally, because participation in both phases was voluntary, participants who were more engaged with or more comfortable discussing their GAI use may be over-represented, a form of self-selection that could inflate the salience of reflective, dialogic accounts relative to less visible or less articulate forms of cognitive delegation.

6. Recommendations

For teaching and assessment. Academics should design AI-permitted tasks that make cognitive responsibility observable. Students can be required to document prompts and revisions, compare generated claims with scholarly sources, identify errors, justify which elements were accepted or rejected, and explain the reasoning that remains their own. Such designs shift assessment from detecting AI use toward evaluating judgement, verification and intellectual ownership.

For institutional governance. Universities should combine clear rules on disclosure, privacy, academic integrity and appropriate delegation with systematic AI-literacy development for students and faculty. Governance should not be reduced to prohibition or detection. The modest moderation effect observed here suggests that institutional guidance can support productive use, but it cannot substitute for learner regulation.

For equity and capacity development. Institutions should treat critical AI use as a capability-development issue. Access initiatives should be accompanied by support for prompting, source verification, metacognitive monitoring and responsible decision making, especially where connectivity, devices and formal preparation are uneven. The proposed notion of cognitive orchestration inequality should be tested rather than assumed, but it provides a useful hypothesis for examining whether unequal regulatory capacity becomes a new layer of digital inequality.

For research. ACA should next be subjected to independent psychometric validation. Priority tests include confirmatory replication of its dimensional structure, test–retest reliability, convergent and discriminant validity, incremental validity beyond self-regulation, critical thinking and AI literacy, and measurement invariance across students and faculty, linguistic groups and institutional sectors. Longitudinal and experimental

studies should examine temporal ordering and causal mechanisms, while trace data and performance tasks should complement self-report. Multilevel and culturally responsive work should also test whether collective verification and peer-supported human–AI orchestration represent forms of autonomy not fully captured by an individually administered scale.

7. Conclusion

This study examined not whether generative AI is intrinsically beneficial or harmful, but under what conditions its pedagogical use is associated with cognitive autonomy and deep learning. Across six Cameroonian higher-education institutions, GAI use was positively associated with self-regulation; self-regulation and cognitive engagement formed a significant sequential indirect pathway to ACA; ACA was strongly associated with deep learning; and AI literacy and institutional governance showed positive but modest moderation effects. The interview evidence qualified these average relationships by revealing both reflective, dialogic use and cognitive delegation. Generative AI does not remove the need for learner autonomy; it changes what autonomy requires. In this study, productive GAI use was associated with self-regulation, cognitive engagement and deep learning when users retained control over verification, strategy and intellectual responsibility. Interviews simultaneously showed that the same tools can encourage delegation and reduced effort. The central educational issue is therefore not simply access or frequency of use, but the quality of human-AI orchestration.

The emerging Augmented Cognitive Autonomy framework offers a way to study that quality without treating technology and autonomy as opposites. For higher education in resource-diverse settings such as Cameroon, the implication is practical: expanding access to generative AI should be accompanied by AI literacy, pedagogical redesign and institutional guidance that protect human judgement. If these capabilities are unevenly distributed, GAI may amplify not only learning but inequality. Learning for development in the AI era consequently depends on building people's capacity to use artificial agents without surrendering the cognitive work through which understanding, agency and responsibility are formed. Whether this capacity can itself be taught, and how quickly institutions in resource-constrained settings can build the guidance it requires, remain the practical questions this framework is intended to help answer.

Declarations

Author contributions: René Blaise Mbop Ackok: Conceptualization, Methodology, Investigation, Formal Analysis, Project Administration, Writing – Original Draft, Writing – Review & Editing.

Funding: The study received no financial support.

Ethical statement: Participation was voluntary and based on informed consent. Confidentiality and anonymisation were maintained throughout data collection, analysis, and reporting. Ethical approval was obtained from the Ethics Committee of the Higher Institute of Educational Sciences, Adventist University Cosendai, Cameroon. All procedures were conducted in accordance with the ethical principles applicable to research involving human participants.

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

Data sharing statement: Data supporting the findings of this study are available from the corresponding author upon reasonable request, subject to ethical and confidentiality constraints.

Use of AI: ChatGPT (OpenAI) was used during manuscript revision to support language refinement, structural editing, and clarity. The tool was not used to generate or alter the study data, conduct the reported statistical analyses, or independently determine the scientific conclusions. The author critically reviewed and validated the revised content and takes full responsibility for the scientific content, analyses, interpretations, and final manuscript.

References

Bandura, A. (1986). *Social foundations of thought and action: A social cognitive theory.* Prentice-Hall.

Biggs, J., & Tang, C. (2011). *Teaching for quality learning at university* (4th ed.). Open University Press.

Boateng, G. O., Neilands, T. B., Frongillo, E. A., Melgar-Quiñonez, H. R., & Young, S. L. (2018). Best practices for developing and validating scales for health, social, and behavioral research: A primer. *Frontiers in Public Health, 6,* 149. https://doi.org/10.3389/fpubh.2018.00149

Braun, V., & Clarke, V. (2022). *Thematic analysis: A practical guide.* SAGE.

Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education, 20,* 43. https://doi.org/10.1186/s41239-023-00411-8

Chiu, T. K. F. (2024). A classification tool to foster self-regulated learning with generative artificial intelligence by applying self-determination theory: A case of ChatGPT. Educational Technology Research and Development, 72, 2401–2416. https://doi.org/10.1007/s11423-024-10366-w

Creswell, J. W., & Plano Clark, V. L. (2018). *Designing and conducting mixed methods research* (3rd ed.). SAGE.

Deci, E. L., & Ryan, R. M. (2000). The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry, 11*(4), 227–268. https://doi.org/10.1207/S15327965PLI1104_01

DeVellis, R. F., & Thorpe, C. T. (2021). *Scale development: Theory and applications* (5th ed.). SAGE.

Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., et al. (2023). "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management, 71,* 102642. https://doi.org/10.1016/j.ijinfomgt.2023.102642

Fredricks, J. A., Blumenfeld, P. C., & Paris, A. H. (2004). School engagement: Potential of the concept, state of the evidence. *Review of Educational Research, 74*(1), 59–109. https://doi.org/10.3102/00346543074001059

Hair, J. F., Hult, G. T. M., Ringle, C. M., & Sarstedt, M. (2022). *A primer on partial least squares structural equation modeling (PLS-SEM)* (3rd ed.). SAGE.

Helsper, E. J. (2021). *The digital disconnect: The social causes and consequences of digital inequalities.* SAGE.

Hutchins, E. (1995). *Cognition in the wild.* MIT Press.

Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., et al. (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

Li, C., Cui, H., & Hagedorn, L. S. (2026). The cognitive impact of ChatGPT in higher education: A systematic review of critical and creative thinking outcomes. Computers and Education: Artificial Intelligence, 10, 100571. https://doi.org/10.1016/j.caeai.2026.100571

Lodge, J. M., Howard, S., Bearman, M., & Dawson, P. (2023). *Assessment reform for the age of artificial intelligence.* Tertiary Education Quality and Standards Agency.

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems,* 1–16. https://doi.org/10.1145/3313831.3376727

Misiejuk, K., López-Pernas, S., Kaliisa, R., & Saqr, M. (2026). Cognitive offloading in student–AI collaboration: A longitudinal analysis of prompting strategies. Computers in Human Behavior Reports, 22, 101130. https://doi.org/10.1016/j.chbr.2026.101130

Moluayonge, G. E. (2025). Cognitive offload as a mediator of AI usage and critical thinking development in Cameroonian higher education learners. Asian Journal of Education and Social Studies, 51(8), 467–476. https://doi.org/10.9734/ajess/2025/v51i82253

Morgan, D. L. (2014). *Integrating qualitative and quantitative methods: A pragmatic approach.* SAGE.

Panadero, E. (2017). A review of self-regulated learning: Six models and four directions for research. *Frontiers in Psychology, 8,* 422. https://doi.org/10.3389/fpsyg.2017.00422

Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A framework for ethical integration of generative AI in educational assessment. *Journal of University Teaching & Learning Practice, 21*(6). https://doi.org/10.53761/q3azde36

Salomon, G. (Ed.). (1993). *Distributed cognitions: Psychological and educational considerations.* Cambridge University Press.

Tate, T. P., Harnick-Shapiro, B., Ritchie, D. R., Tseng, W., Dennin, M., & Warschauer, M. (2025). Incorporating generative AI into a writing-intensive undergraduate course without off-loading learning. Discover Computing, 28, Article 72. https://doi.org/10.1007/s10791-025-09563-9

Tlili, A., Shehata, B., Adarkwah, M. A., Bozkurt, A., Hickey, D. T., Huang, R., & Agyemang, B. (2023). What if the devil is my guardian angel? ChatGPT as a case study of using chatbots in education. *Smart Learning Environments, 10,* 15. https://doi.org/10.1186/s40561-023-00237-x

UNESCO. (2023). *Guidance for generative AI in education and research.* UNESCO.

van Dijk, J. (2020). *The digital divide.* Polity Press.

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes.* Harvard University Press.

Wu, X.-Y., & Chiu, T. K. F. (2025). Integrating learner characteristics and generative AI affordances to enhance self-regulated learning: A configurational analysis. Journal of New Approaches in Educational Research, 14, Article 10. https://doi.org/10.1007/s44322-025-00028-x

Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice, 41*(2), 64–70. https://doi.org/10.1207/S15430421TIP4102_2