FROM CONTENT TRANSMISSION TO EPISTEMIC CONSTRUCTION: REDESIGNING UNIVERSITY EDUCATION IN THE AGE OF GENERATIVE ARTIFICIAL INTELLIGENCE

Ronald Jean Degen

ANEGEPE - Academia de Estudos em Empreendedorismo e Gestão de Pequenas Empresas São Paulo, Brazil

ABSTRACT

Generative artificial intelligence challenges the university model centered on content transmission and subsequent testing. Large language models can retrieve, synthesize, and recombine information within seconds, but rapid access to information does not constitute reliable knowledge. This conceptual article argues that higher education should shift toward epistemic construction: the process of framing problems, integrating information, verifying sources, exposing assumptions, confronting counterarguments, revising provisional judgments, and defending conclusions. Drawing on Socratic inquiry, medieval disputatio, Popperian falsification, reflective learning, critical-thinking research, and scholarship on artificial intelligence in education, the study proposes the AI-Assisted Epistemic Construction Cycle. The framework positions generative AI not primarily as an answer-producing device, but as a scalable digital devil's advocate that supports structured intellectual dissent. Implications are developed for faculty roles, student agency, assessment, institutional governance, academic integrity, and future research.

KEYWORDS

generative artificial intelligence, critical thinking, epistemic construction, higher education, structured dissent

1. INTRODUCTION

For much of their history, universities operated under conditions of relative informational scarcity. Books were expensive, libraries were geographically concentrated, and access to disciplinary expertise depended heavily on institutional membership. Under these conditions, the lecture-based model served an essential function: faculty selected, organized, interpreted, and transmitted knowledge that students could not easily obtain elsewhere. This historical context produced relatively linear educational architecture. Instructors presented disciplinary content; students studied and practiced it; and institutions evaluated whether students could reproduce, explain, or apply what they had received. Lectures, textbooks, examinations, and essays became the principal mechanisms through which knowledge was distributed and learning was certified.

Generative artificial intelligence fundamentally alters the informational environment in which this architecture operates. Large language models can produce explanations, summaries, comparisons, case analyses, research outlines, computer code, and polished prose almost instantaneously. Their educational use presents substantial opportunities, including individualized explanation and rapid formative feedback, but also risks involving hallucinations, bias, overreliance, privacy, and academic integrity [1], [2].

The resulting challenge is not simply technological. It is epistemological and pedagogical. If sophisticated answers can be generated at negligible marginal cost, possession of a polished answer can no longer be treated as sufficient evidence that learning has occurred. A student may produce a coherent essay without understanding the concepts, evidence, or assumptions embedded in it. Performance and learning can increasingly diverge.

Universities therefore face a fundamental question: What remains their distinctive educational contribution when information and plausible answers are continuously available?

This article argues that universities should not attempt to compete with generative AI in the rapid transmission of information. Nor should they respond primarily by prohibiting the technology. Instead, they should shift the center of education from content transmission to epistemic construction: the disciplined process through which learners transform information into knowledge that is conceptually integrated, evidentially supported, critically tested, contextually interpreted, and open to revision.

This shift does not imply that foundational knowledge, memory, lectures, or disciplinary expertise have become unnecessary. Critical judgment cannot operate without substantive knowledge. Students who lack conceptual foundations cannot reliably detect errors, evaluate sources, or recognize implausible claims. The proposed transformation is therefore not a movement from knowledge to generic skills, but from passive reception of knowledge claims to active construction and critical justification of knowledge.

2. METHOD

This study adopts a conceptual theory-synthesis approach. Following Jaakkola [3], its purpose is not to test a predefined causal model or provide an exhaustive systematic review, but to reorganize established concepts from different research traditions in order to develop a theoretically coherent explanation and a pedagogical framework. The analysis integrates four bodies of literature that are frequently discussed separately: critical-thinking and active-learning research [4]; historical practices of institutionalized intellectual challenge; experiential, reflective, and organizational learning; and contemporary scholarship on generative artificial intelligence in education. Seminal books, peer-reviewed articles, and authoritative institutional guidance were purposively selected for their relevance to the central research question: how should university education be redesigned when students can obtain plausible answers and synthesized information almost instantaneously?

The conceptual procedure involved four stages:

First, the transmissive model was reconstructed in relation to the conditions of informational scarcity under which it developed.

Second, information access was analytically distinguished from knowledge construction and judgment.

Third, common epistemic principles were identified across Socratic inquiry, medieval disputatio, falsification, peer review, and reflective learning.

Fourth, these principles were synthesized with learner-centered uses of generative AI to formulate the AI-Assisted Epistemic Construction Cycle and a set of research propositions.

The resulting framework is therefore an original integrative synthesis rather than a reproduction of a previously published learning cycle. Its credibility rests on conceptual coherence, explicit links to the literature, transparency regarding its non-empirical status, and the formulation of propositions that can be evaluated in future research.

3. RESULTS

The conceptual synthesis produced three interrelated results:

First, it clarifies why informational abundance weakens the validity of educational models centered on content reproduction.

Second, it identifies structured critical challenge as a recurrent epistemic mechanism across historical and contemporary traditions.

Third, it translates these insights into the AI-Assisted Epistemic Construction Cycle, an eight-stage framework for teaching, learning, and assessment.

3.1. The Limits of Content Transmission in an Age of Information Abundance

The transmissive model rests on several implicit assumptions. The first is that instructors possess knowledge that students cannot readily access without expert mediation. The second is that well-organized exposure to this knowledge produces learning. The third is that reproduction or application of transmitted content provides valid evidence of understanding.

These assumptions were always incomplete. Freire's [5] critique of banking education challenged the idea that knowledge could simply be deposited by teachers into passive students. Constructivist and active-learning perspectives similarly emphasized that learners do not merely receive meaning; they construct it through engagement with concepts, experiences, problems, and other people.

Research on cognitive engagement supports this distinction. The ICAP (Interactive, Constructive, Active, and Passive) framework differentiates passive, active, constructive, and interactive modes of learning. Its central proposition is that learning deepens when students move from receiving information to manipulating it, generating inferences, constructing explanations, and interacting with competing ideas [6].

Generative AI makes the weaknesses of purely transmissive education more visible. A polished essay, correct definition, or well-organized presentation can no longer be assumed to represent the student's own conceptual work. The final product may conceal substantial cognitive outsourcing. This problem existed before generative AI, but the technology greatly increases its scale and lowers its cost.

This does not mean that lectures, textbooks, or examinations should disappear. Expert explanations can provide conceptual structures that novices would struggle to construct independently. Foundational knowledge must often be internalized before higher-order evaluation becomes possible. A medical student requires knowledge of anatomy and physiology; an engineer requires mathematical and physical principles; a lawyer requires legal concepts and precedents; and a manager requires knowledge of finance, organizations, strategy, and markets.

The necessary transformation is therefore not the elimination of content but its repositioning.

Disciplinary knowledge should become the material through which students learn to interpret, compare, question, integrate, and judge. Content remains necessary, but its reproduction can no longer serve as the dominant definition of learning.

3.2. Access to Information is Not Knowledge

Claims that AI gives students immediate access to all knowledge require conceptual qualification. Large language models provide rapid access to linguistic representations generated from patterns in human-produced data. Their output may be accurate, incomplete, biased, outdated, or fabricated. Fluency is not equivalent to truth, and grammatical coherence is not equivalent to understanding [7].

A useful distinction can be drawn among data, information, knowledge, and judgment. Data consists of observations, measurements, signs, or records. Information emerges when data is organized and interpreted. Knowledge requires the integration of information into conceptual structures that allow explanation, comparison, inference, and application. Judgment concerns the selection of what is relevant, credible, ethically acceptable, and appropriate in a specific context.

Generative AI can assist at each level, but it does not remove the need for human epistemic responsibility. A student may obtain an explanation of inflation, genetic inheritance, constitutional interpretation, or competitive strategy without acquiring the ability to determine whether that explanation is accurate or applicable.

External access can also create an illusion of internal understanding. Fisher et al. [8] found that searching for explanations online can inflate individuals' estimates of their own knowledge. Sparrow et al. [9] similarly demonstrated that when people expect information to remain externally available, they may remember where to find it more readily than the information itself.

Cognitive offloading is not inherently harmful. Writing, libraries, calculators, databases, and search engines have long extended human cognition. The risk arises when external access replaces the internal conceptual structures required to interpret what has been retrieved. Critical thinking therefore cannot be taught independently of disciplinary knowledge. A learner must know enough to recognize an anachronism, detect an implausible causal claim, distinguish a valid inference from a rhetorical assertion, and identify when a source has been misrepresented.

The educational goal should not be the accumulation of every available fragment of information. It should be the development of epistemic vigilance: the ability to assess the plausibility of a claim, the quality of its evidence, and the reliability of its source [10].

Epistemic construction consequently involves defining what is being claimed, locating the claim within a conceptual and historical context, identifying the evidence on which it depends, examining assumptions and alternative explanations, determining the conditions under which it is valid, and remaining willing to revise it when stronger evidence or arguments emerge. Under this conception, knowledge is neither a mere collection of information nor an immutable possession. It is a justified and revisable achievement.

3.3. Institutionalized Critical Challenge as an Epistemic Tradition

The proposal to organize learning around questioning and criticism is not a technological invention. It belongs to a long intellectual history in which communities developed procedures for exposing claims to disciplined challenge before accepting them as knowledge.

Socratic Inquiry

In the Socratic tradition, learning did not occur primarily through the transmission of completed doctrines. An interlocutor was invited to define a concept or defend a belief. Subsequent questioning exposed ambiguity, contradiction, or consequences that had not been considered.

The Socratic elenchus shifted attention from the conclusion to its justification. It required participants to articulate what they believed, why they believed it, and whether their beliefs remained coherent when examined alongside other commitments [11].

The recognition of ignorance was not treated as educational failure. It was the beginning of inquiry. This insight remains central in the age of AI: learners do not achieve understanding merely by receiving an answer; they develop understanding by explaining, testing, and revising the reasons for accepting it.

Medieval Disputatio

Medieval universities institutionalized a more formal version of intellectual challenge through the disputatio. A question was proposed, arguments on opposing sides were examined, objections were addressed, and a reasoned determination was formulated.

The disputatio was not simply unregulated debate. It was a structured pedagogical and scholarly practice governed by conventions of argument and response [12]. Its ideal purpose was not to humiliate an opponent but to refine an idea through criticism.

This practice established an important institutional principle: disagreement can be designed as educational infrastructure. When criticism is governed by shared norms, directed toward claims rather than persons, and followed by reasoned reconstruction, it improves the quality of judgment.

Science, PeerReview and Reflective Learning

Modern science extended the same epistemic logic. Popper [13] argued that scientific progress depends not merely on accumulating confirming observations, but on exposing hypotheses to serious attempts at refutation.

Peer review similarly institutionalizes external scrutiny. Although imperfect, it embodies the recognition that authors are not always the best judges of their own assumptions, methods, and conclusions.

Reflective-learning theories add a developmental dimension. Kolb [14] conceptualized learning as a cycle involving experience, reflection, conceptualization, and experimentation. Schön [15] emphasized reflection within and upon professional practice. Argyris and Schön [16] distinguished the correction of actions from the deeper examination of the assumptions governing those actions.

These traditions converge on a common proposition: knowledge and judgment improve when provisional interpretations encounter structured challenge, reflection, revision, and reapplication.

Generative AI as Source and Digital Devil's Advocate

Most educational discussions initially treated generative AI as a tool for producing answers.

Students could request explanations, summaries, essays, code, or solutions. This use may improve efficiency, but it can also reproduce the passivity of traditional instruction: the human lecturer is replaced by a machine lecturer, while the learner remains principally a recipient.

The more transformative educational role of generative AI may lie in its ability to act as a digital devil's advocate. In this role, AI is not asked primarily to complete the task. It is instructed to expose the learner's reasoning to organized criticism.

A system can be asked to identify hidden assumptions, distinguish evidence from inference, formulate the strongest counterargument, generate counterexamples, search for disconfirming evidence, represent competing theoretical traditions, simulate the perspectives of affected stakeholders, identify unintended consequences, specify the conditions under which a conclusion would fail, or conduct a sequence of Socratic questions.

This use of AI does not make the system an epistemic authority. Its objections may themselves be superficial, biased, or false. The learner must evaluate the criticism rather than submit to it. The educational relationship is therefore recursive: the student questions the AI; the AI questions the student; and the resulting claims are checked against scholarly sources, empirical evidence, peers, and faculty expertise.

Ouyang and Jiao's [17] distinction among AI-directed, AI-supported, and AI-empowered education is relevant here. The digital devil's advocate belongs to a learner-centered model in which AI supports but does not displace human agency.

The educational significance of this function is scalability. Historically, sustained intellectual challenge depends on access to skilled instructors, expert reviewers, doctoral seminars, mentors, or highly capable peers. Generative AI can make some forms of structured questioning continuously available. Its value lies not in guaranteeing better answers but in increasing the frequency with which provisional answers are subjected to criticism.

Illustrative prompts include:

- Identify every assumption that must be true for this conclusion to remain valid.
- Present the strongest argument that an informed critic could make against my position.
- What evidence would falsify this claim?
- Which affected stakeholder is absent from my analysis?
- Do not agree with me automatically. Ask one Socratic question at a time until the weaknesses of my argument are explicit.

3.4. The AI-Assisted Epistemic Construction Cycle

To operationalize the transition from transmission to epistemic construction, this article proposes the AI-Assisted Epistemic Construction Cycle (AECC). The framework is an integrative conceptual synthesis rather than a reproduction of a previously established learning cycle. It draws on Socratic inquiry, medieval disputatio, Popperian falsification, experiential and reflective learning, critical-thinking research, and learner-centered approaches to artificial intelligence.

The AECC comprises eight recursive stages. It does not seek to eliminate direct instruction. Instead, it places instruction within a broader architecture in which provisional judgments are repeatedly informed, tested, reconstructed, defended, and reapplied (Figure 1):

Stage 1: Problem formulation. Learning begins with a problem requiring interpretation rather

than simple retrieval. The problem may contain incomplete information, competing values, uncertain causal relationships, or more than one defensible response. Students must determine what the problem is, who is affected, which variables matter, and what remains unknown. Poor answers often originate in poorly framed questions.

Stage 2: Initial judgment without AI. Before consulting AI, students formulate and record a provisional interpretation. They identify the knowledge they are using, the assumptions they are making, and the uncertainties they recognize. This stage preserves intellectual agency and reduces anchoring on the first AI-generated response. It also creates a baseline against which the evolution of reasoning can be assessed.

Stage 3: Mapping relevant information with AI. Students then use academic databases, disciplinary sources, and, where appropriate, generative AI to map the relevant knowledge landscape. The purpose is not to collect the largest possible quantity of material, but to distinguish among established findings, contested interpretations, theoretical claims, empirical evidence, contextual variables, and unresolved questions.

Stage 4: Source verification. Every consequential claim must be traced to an identifiable source and evaluated. Students examine whether the source exists, whether it has been peer reviewed, whether its method supports its conclusion, whether it is current, whether contrary evidence exists, and whether the AI represented it accurately. This stage converts the limitations of AI into opportunities for developing epistemic vigilance.

Stage 5: AI-assisted counterargument. Students instruct the AI to act adversarial. It identifies unsupported assumptions, generates objections, proposes alternative explanations, and searches for circumstances in which the argument would fail. Students must classify each objection as relevant, partially relevant, irrelevant, or mistaken, explaining the basis for their judgment.

Stage 6: Human disputatio. AI-supported criticism is complemented by human dialogue. Students defend competing positions, question one another, and respond to arguments that may reflect experience, values, tacit knowledge, or contextual understanding unavailable to the system. Where appropriate, learners should be asked to defend a position different from their initial preference. This reduces personal identification with a conclusion and encourages charitable reconstruction of opposing arguments.

Stage 7: Reflective reconstruction. Students revise their original judgment. They document which assumptions were retained, modified, or rejected; which evidence changed in importance; which objections were incorporated; and which uncertainties remain. Revision should be treated as evidence of learning, not inconsistency. Intellectual maturity includes the willingness to change one's mind for defensible reasons.

Stage 8: Defense and application. The reconstructed conclusion is defended orally or in writing and then applied to a new context. Faculty may introduce new evidence, alter the conditions of the problem, or present a counterexample. Transfer to a different case test whether students understand the underlying principles or merely reproduce the previous solution. Application produces experience and feedback, which may require renewed problem formulation and a revised judgment.

The cycle therefore remains recursive.

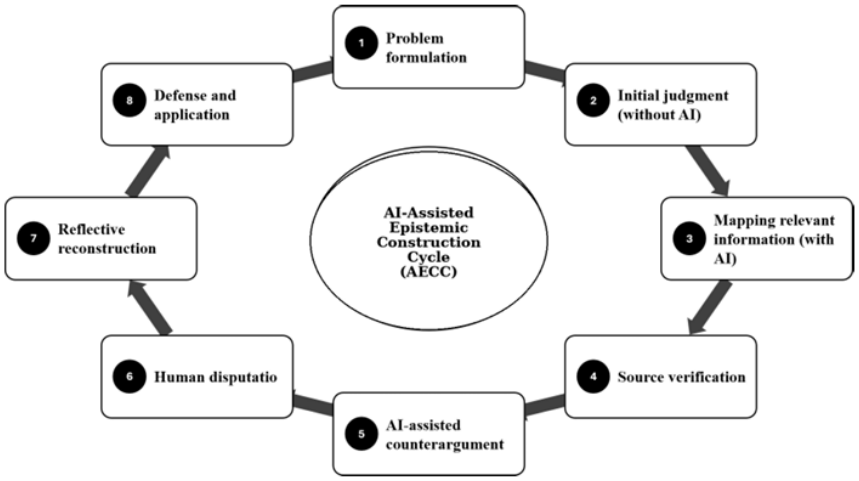

Figure 1. AI-Assisted Epistemic Construction Cycle (AECC).

Source: Elaborated by the author from Socratic inquiry [11], medieval disputatio [12], Popperian criticism [13], experiential and reflective learning [14]-[16], epistemic vigilance [10], and learner-centered AI [1], [17].

4. DISCUSSION

The AECC reframes the educational implications of generative AI. Its central contribution is not the addition of a technological tool to an unchanged curriculum, but a change in the unit of pedagogical design: from delivering completed answers to constructing learning architectures in which provisional judgments encounter information, verification, criticism, revision, and application.

4.1. Redefining the Roles of Faculty and Students

The proposed model does not diminish the importance of faculty. It makes their role more intellectually demanding. Faculty moves from being principally transmitters of information to becoming designers of epistemic environments. They select problems, sequence inquiries, identify foundational knowledge, establish standards of evidence, and create psychologically safe conditions for disagreement.

Faculty also acts as disciplinary curators. Novices cannot always distinguish central theories from marginal claims or rigorous evidence from plausible misinformation. Expert judgment remains essential.

Instructors additionally become moderators of structured dissent. Productive disagreement requires norms that separate criticism of ideas from attacks on people. The instructor ensures that objections are relevant, arguments are represented fairly, and disagreement leads to reconstruction rather than polarization.

Finally, faculty evaluates the quality of reasoning rather than only the correctness of the final answer. Students cease to be passive recipients and assume responsibility as problem framers, source auditors, argument constructors, and revisers of provisional judgments. They must learn the difference between using AI as a cognitive instrument and delegating to AI the responsibility to think. Generative systems may suggest, summarize, compare, criticize, and reorganize. They

cannot assume moral, professional, or legal responsibility for the consequences of a decision. Accountability remains human.

4.2. Assessment of Epistemic Construction

Assessment must be redesigned if universities intend to value epistemic construction. Biggs's [18] principle of constructive alignment requires coherence among learning objectives, teaching activities, and assessment. If a program declares that critical thinking is central but rewards primarily content reproduction, its actual curriculum remains transmissive.

Assessment within the AECC should combine evidence of initial thinking, source use, criticism, revision, transfer, and unaided competence. The aim is to make the development of reasoning visible rather than infer it solely from a final product.

Rubrics should evaluate problem formulation, conceptual precision, source quality, recognition of assumptions, treatment of counterevidence, coherence between evidence and conclusion, willingness to revise, transfer, and transparency in AI use.

Abrami et al.'s [19] meta-analysis found that dialogue, authentic or situated problems, and mentoring contribute positively to the development of critical thinking. The AECC embeds these conditions in a recursive architecture.

Assessment should also preserve occasions for unaided performance. Certain knowledge and capabilities must be available without technological assistance, especially in professional contexts where immediate judgment is required or AI access is inappropriate. The goal is neither total prohibition nor unrestricted use. It is calibrated delegation.

Accordingly, an aligned assessment system may combine an unaided initial judgment, an annotated map of sources and claims, a record of significant AI interactions, a counterargument analysis, a revision memorandum, a human disputatio, oral defense, transfer to a new case, and selected tasks completed without AI.

4.3. Institutional Governance, Academic Integrity and Transparency

The integration of AI requires institutional governance rather than isolated classroom rules. Universities should distinguish among activities in which AI is prohibited, permitted, encouraged, or required. The appropriate rule depends on the learning objective. An examination of foundational knowledge may prohibit AI, whereas an assignment on source verification may require its use.

Transparency should become a normal feature of academic work. Students and researchers should disclose the system used, the purpose for which it was used, the stages in which it contributed, the procedures used to verify its outputs, and the substantive suggestions accepted, modified, or rejected.

Disclosure should not automatically be treated as evidence of misconduct. It can instead demonstrate methodological transparency and epistemic responsibility. At the same time, disclosure alone is insufficient. Understanding should be tested through oral defense, transfer, source verification, and comparison between initial and revised reasoning.

Institutions must also address privacy, confidential data, intellectual property, unequal access to advanced systems, and the risk that commercially available tools may influence curricular

priorities. UNESCO's guidance emphasizes human agency, inclusion, privacy, transparency, and validation as central conditions for responsible educational adoption [2].

Zawacki-Richter et al. [20] found that much of the earlier AI-in-higher-education literature was driven more strongly by technical development than by educators and pedagogical theory. The lesson remains relevant: educational purposes should govern technology, not the reverse.

4.4. Risks and Safeguards

A model centered on AI-assisted epistemic construction must recognize that the technology can weaken learning as well as support it. The principal risks require explicit pedagogical and institutional safeguards:

- Cognitive outsourcing. Students may transfer problem formulation, reasoning, and evaluation to the system. Requiring an unaided initial judgment and selected no-AI assessments preserves foundational capability.
- Illusions of understanding. Fluent outputs may be mistaken for comprehension. Oral defense, transfer tasks, and explanation in the student's own words help reveal whether conceptual understanding exists. • Fabricated or distorted sources. Every consequential citation and factual claim must be checked against the original source.
- Confirmation and accommodation. Systems may align too readily with a user's position. Prompts should explicitly require disconfirming evidence and the strongest opposing interpretation.
- Bias and exclusion. Students should examine whose perspectives, interests, languages, and experiences are absent from the generated response.
- Privacy and confidentiality. Personal data, unpublished research, proprietary information, and sensitive professional material should not be uploaded without appropriate safeguards.
- Inequality of access. Universities should avoid designing assessments in which access to expensive systems becomes an unacknowledged academic advantage.
- Erosion of human dialogue. AI-supported criticism should complement rather than replace interaction with faculty, peers, practitioners, and affected communities.

4.5. Ibero-American Higher Education

Ibero-American institutions face shared opportunities but unequal conditions. Research from the region reports interest in AI as a learning mediator and a need to move beyond simple prohibition [21]-[23]. Recent evidence also indicates that GenAI can support formative feedback when embedded in a pedagogical protocol rather than used as an autonomous evaluator [24].

Implementation must nevertheless account for unequal connectivity, access to paid models, faculty workload, institutional capacity, accessibility, and linguistic asymmetry. Models trained predominantly on English material may underrepresent Portuguese, Spanish, Indigenous, and locally situated knowledge. Public universities cannot assume that every student has equivalent devices, subscriptions, or private study conditions. Requiring a premium system without institutional provision would convert financial capacity into an unacknowledged academic advantage.

Regional adoption should therefore prioritize multilingual verification, equitable access, faculty development, data protection, and locally governed policies. Governance studies emphasize that effective adoption requires coordinated frameworks, stakeholder participation, and clear responsibilities rather than fragmented classroom rules [25]. The AECC is best treated as an

adaptable architecture: disciplinary communities should define acceptable evidence, unaided competencies, ethical boundaries, and culturally appropriate forms of disagreement.

5. THEORETICAL CONTRIBUTION AND RESEARCH AGENDA

The framework generates propositions that can be examined empirically across disciplines, educational levels, and cultural contexts:

Proposition 1. Students who use generative AI primarily to challenge assumptions will develop stronger critical judgment than students who use it primarily to generate completed answers.

Proposition 2. Repeated exposure to structured intellectual dissent will produce greater gains in epistemic vigilance than equivalent exposure to unchallenged content.

Proposition 3. Requiring an unaided initial judgment before AI use will reduce anchoring and increase learners' ability to identify changes in their reasoning.

Proposition 4. The benefits of AI-assisted counterargument will be moderated by disciplinary knowledge; students with stronger conceptual foundations will evaluate AI criticism more effectively.

Proposition 5. Combining AI-assisted counterargument with human disputatio will produce deeper reflective reconstruction than either mechanism alone.

Proposition 6. Assessment systems that document revision and require oral defense will provide more valid evidence of understanding than final written products considered in isolation.

Proposition 7. Educational environments that reward justified revision will develop greater intellectual humility than environments that reward persistent defense of an initial answer.

Proposition 8. AI integration without explicit instruction in source verification and epistemic vigilance will increase the risk of confident but poorly justified conclusions.

Future research should test the AECC across disciplines, educational levels, cultures, and technological systems. Studies should compare different forms of AI prompting, determine the optimal frequency of adversarial interaction, examine effects on long-term retention and transfer, and investigate whether intensive criticism may create cognitive overload or reduce confidence among less-prepared learners.

6. INTERPRETATION AND LIMITATIONS

The emergence of generative AI does not eliminate the educational need for knowledge. It clarifies the difference between possessing information and being capable of judging it.

The traditional transmission model was appropriate to an environment in which access to organized knowledge was limited. In an environment of informational abundance, the university's distinctive contribution shifts toward selection, verification, integration, criticism, and responsible application.

This shift also reframes critical thinking. It should not be treated solely as an individual trait that some students possess more strongly than others. Critical thinking can be institutionally produced

through roles, routines, assignments, and assessment systems that require ideas to encounter disciplined opposition.

The AECC therefore changes the unit of analysis from the isolated critical thinker to the learning architecture. It asks not only whether students are disposed to question, but whether the institution makes questioning, source verification, counterargument, and revision unavoidable components of learning.

Generative AI is relevant because it can expand access to structured challenges. It does not inaugurate critical thinking and should not be anthropomorphized as a conscious Socratic partner. Its contribution is instrumental: it can generate questions, objections, comparisons, and alternative perspectives at scale.

The final judgment, however, must remain human. AI can identify a possible contradiction, but humans must determine whether it is genuine. It can suggest a reference, but researchers must verify that the reference exists and supports the claim. It can recommend a decision, but it cannot bear responsibility for its consequences.

The article has four principal limitations:

First, the AECC is a conceptual model and has not yet been empirically validated.

Second, the framework is intentionally general and may require adaptation to disciplinary epistemologies, professional regulations, student age, and institutional resources.

Third, the effectiveness of AI-assisted counterargument will depend on model quality, prompt design, language, and access conditions.

Fourth, structured criticism may generate cognitive overload or reduce confidence when introduced without adequate scaffolding and psychological safety.

These limitations define a research agenda rather than invalidate the framework. Comparative studies should examine different disciplines, levels of prior knowledge, forms of AI prompting, frequencies of adversarial interaction, and combinations of machine-generated and human criticism. Longitudinal research is particularly necessary to determine whether epistemic construction improves retention, transfer, intellectual humility, and professional judgment.

7. CONCLUSION

Generative artificial intelligence has sharply reduced the cost of producing sophisticated answers. As a result, possession of an answer can no longer serve as reliable evidence of learning.

Universities should respond by moving beyond an educational model centered predominantly on content transmission and reproduction. Their central mission should increasingly become the cultivation of epistemic construction: the ability to transform information into knowledge that is integrated, verified, challenged, justified, and responsibly applied.

This transformation does not require abandoning disciplinary knowledge, lectures, or human instruction. On the contrary, it increases the importance of expert faculty and foundational knowledge. Students cannot critically evaluate what they do not understand.

The Socratic tradition demonstrated the educational power of questioning. Medieval disputatio institutionalized the examination of opposing arguments. Popperian epistemology emphasized the value of attempted refutation. Reflective-learning traditions showed how judgment develops through experience, criticism, reconstruction, and application.

Generative AI can extend this historical trajectory when it is used not merely to generate answers but to institutionalize critical challenges. Configured as a digital devil's advocate, it can expose assumptions, formulate counterarguments, simulate alternative perspectives, and prompt reflective revision.

The distinctive role of university in the age of AI is therefore not to compete with machines in producing information. It is to develop the human capacity to judge information. The essential educational question is no longer simple, what does the student know? It is also: How was the conclusion constructed? Which sources were verified? Which assumptions were identified? Which objections were confronted? What evidence would justify revision?

The future of higher education lies in this movement from information to understanding, from answers to justification, and from transmission to epistemic construction.

8. STATEMENT ON THE USE OF ARTIFICIAL INTELLIGENCE

Generative artificial intelligence was used during the development of this article to support translation, manuscript organization, argument refinement, language revision, and the generation of counterarguments. It was also used as a digital devil's advocate to identify assumptions, possible inconsistencies, and alternative interpretations.

All AI-generated contributions were critically reviewed, verified, modified, or rejected by the author. The author assumes full responsibility for the theoretical propositions, interpretations, reference selection, argumentation, academic integrity, and final content. Artificial intelligence was not treated as an author or coauthor because it cannot assume intellectual, ethical, or legal responsibility for the work.

REFERENCES

[1] E. Kasneci et al., “ChatGPT for good? On opportunities and challenges of large language models for education,” Learning and Individual Differences, vol. 103, Art. no. 102274, 2023, doi: 10.1016/j.lindif.2023.102274.

[2] F. Miao and W. Holmes, Guidance for Generative AI in Education and Research. Paris, France: UNESCO, 2023, doi: 10.54675/EWZM9535.

[3] E. Jaakkola, “Designing conceptual articles: Four approaches,” AMS Review, vol. 10, pp. 18-26, 2020, doi: 10.1007/s13162-020-00161-0.

[4] P. A. Facione, Critical Thinking: A Statement of Expert Consensus for Purposes of Educational Assessment and Instruction—The Delphi Report. Millbrae, CA, USA: California Academic Press, 1990.

[5] P. Freire, Pedagogy of the Oppressed, 30th anniversary ed. New York, NY, USA: Continuum, 2000. Original work published 1970.

[6] M. T. H. Chi and R. Wylie, “The ICAP framework: Linking cognitive engagement to active learning outcomes,” Educational Psychologist, vol. 49, no. 4, pp. 219-243, 2014, doi: 10.1080/00461520.2014.965823.

[7] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, “On the dangers of stochastic parrots: Can language models be too big?” in Proc. 2021 ACM Conf. Fairness, Accountability, and Transparency, New York, NY, USA: ACM, 2021, pp. 610-623, doi: 10.1145/3442188.3445922.

[8] M. Fisher, M. K. Goddu, and F. C. Keil, “Searching for explanations: How the Internet inflates

estimates of internal knowledge," Journal of Experimental Psychology: General, vol. 144, no. 3, pp. 674-687, 2015, doi: 10.1037/a0039650.
[9] B. Sparrow, J. Liu, and D. M. Wegner, "Google effects on memory: Cognitive consequences of having information at our fingertips," Science, vol. 333, no. 6043, pp. 776-778, 2011, doi: 10.1126/science.1207745.
[10] D. Sperber et al., "Epistemic vigilance," Mind & Language, vol. 25, no. 4, pp. 359-393, 2010, doi: 10.1111/j.1468-0017.2010.01394.x.
[11] G. Vlastos, Socratic Studies. Cambridge, U.K.: Cambridge University Press, 1994.
[12] A. J. Novikoff, The Medieval Culture of Disputation: Pedagogy, Practice, and Performance. Philadelphia, PA, USA: University of Pennsylvania Press, 2013.
[13] K. R. Popper, The Logic of Scientific Discovery. London, U.K.: Routledge, 2002. Original work published 1959.
[14] D. A. Kolb, Experiential Learning: Experience as the Source of Learning and Development. Englewood Cliffs, NJ, USA: Prentice Hall, 1984.
[15] D. A. Schön, The Reflective Practitioner: How Professionals Think in Action. New York, NY, USA: Basic Books, 1983.
[16] C. Argyris and D. A. Schön, Organizational Learning: A Theory of Action Perspective. Reading, MA, USA: Addison-Wesley, 1978.
[17] F. Ouyang and P. Jiao, "Artificial intelligence in education: The three paradigms," Computers and Education: Artificial Intelligence, vol. 2, Art. no. 100020, 2021, doi: 10.1016/j.caeai.2021.100020.
[18] J. Biggs, "Enhancing teaching through constructive alignment," Higher Education, vol. 32, no. 3, pp. 347 364, 1996, doi: 10.1007/BF00138871.
[19] P. C. Abrami et al., "Strategies for teaching students to think critically: A meta-analysis," Review of Educational Research, vol. 85, no. 2, pp. 275-314, 2015, doi: 10.3102/0034654314551063.
[20] O. Zawacki-Richter, V. I. Marín, M. Bond, and F. Gouverneur, "Systematic review of research on artificial intelligence applications in higher education: Where are the educators?" International Journal of Educational Technology in Higher Education, vol. 16, Art. no. 39, 2019, doi: 10.1186/s41239-019-0171-0.
[21] A. D. Ojeda, A. Solano Barliza, D. Ortega Álvarez, and E. Boom Cárcamo, "Análisis del impacto de la inteligencia artificial ChatGPT en los procesos de enseñanza y aprendizaje en la educación universitaria," Formación Universitaria, vol. 16, no. 6, pp. 61-70, 2023, doi: 10.4067/S0718-50062023000600061.
[22] L. de J. Infante Rivera, M. N. Castillo Rodríguez, G. F. Meza Terbullino, and F. V. Sinche Crispin, "El uso de la inteligencia artificial y su impacto en el aprendizaje de los estudiantes universitarios: Una revisión de la literatura," Revista Ibero-Americana de Estudos em Educação, vol. 19, Art. e18712, 2024.
[23] F. J. García-Peñalvo, "Generative artificial intelligence and education: An analysis from multiple perspectives," Education in the Knowledge Society, vol. 25, Art. e31942, 2024, doi: 10.14201/eks.31942.
[24] M. V. Leiva-Guerrero, I. Araya Zamorano, R. Escobar Collins, and F. Silva Castro, "Retroalimentación de aprendizajes con inteligencia artificial generativa en estudiantes universitarios," RIED—Revista Iberoamericana de Educación a Distancia, vol. 29, no. 1, pp. 241-267, 2026, doi: 10.5944/ried.45547.
[25] M. G. S. Gomes, A. S. M. Ramos, and K. Oliveira, "Guidelines and frameworks of governance for the adoption of generative artificial intelligence in higher education institutions: A scoping review," Revista de Administração da UFSM, vol. 18, no. 4, Art. e1, 2025, doi: 10.5902/1983465990871.