

Artículos atravesados por (o cuestionando) la idea del sujeto -y su género- como una construcción psicobiológica de la cultura.
Articles driven by (or questioning) the idea of the subject -and their gender- as a cultural psychobiological construction.
Vol. 11 (2026), enero-diciembre
ISSN 2469-0783

WHEN AI BECOMES AN INTELLECTUAL INTERMEDIARY: A SITUATED CORRESPONDENCE FRAMEWORK LINKING EPISTEMIC RESPONSIBILITY, INTELLECTUAL ENGAGEMENT, AND INDEPENDENT CAPABILITY IN HIGHER EDUCATION

CUANDO LA IA SE CONVIERTE EN INTERMEDIARIA INTELECTUAL: UN MARCO DE CORRESPONDENCIA SITUADA QUE VINCULA RESPONSABILIDAD EPISTÉMICA, INVOLUCRAMIENTO INTELECTUAL Y CAPACIDAD INDEPENDIENTE EN LA EDUCACIÓN SUPERIOR

Mario Guadalupe López Ayala mario.lopez@uadeo.mx
Universidad Autónoma de Occidente, México.
ORCID: <https://orcid.org/0000-0001-5377-6257>

Israel Osuna Flores iosuna@uaim.edu.mx
Universidad Autónoma Indígena de México, México.
ORCID: <https://orcid.org/0009-0001-0199-2983>

Cómo citar este artículo / Citation: López Ayala, M. G. & Osuna Flores, I. (2026). When AI becomes an intellectual intermediary: A situated correspondence framework linking epistemic responsibility, intellectual engagement, and independent capability in higher education. *Revista Científica Arbitrada de la Fundación MenteClara*, Vol. 11 (443).
DOI: <https://doi.org/10.32351/rca.v11.443>

Copyright: © 2026 RCAFMC. Este artículo de acceso abierto es distribuido bajo los términos de la licencia [Creative Commons Attribution 4.0 International License \(CC BY 4.0\)](https://creativecommons.org/licenses/by/4.0/).
Recibido: 23/09/2026. Aceptado: 28/09/2026. Publicación online: 02/10/2026.

Conflicto de intereses: Ninguno que declarar.

Abstract

Generative artificial intelligence has unsettled a familiar educational inference: that the quality of an academic product is adequate evidence of the intellectual work performed by the student who submits it. This conceptual article develops the Situated Correspondence Framework, an analytical architecture that connects three moments often treated separately: the intellectual operations a task assigns to the learner, evidence that those operations were enacted during AI-mediated activity, and

the capabilities that remain demonstrable when assistance conditions change. The framework distinguishes Assigned Epistemic Responsibility (AER) and Constitutive Intellectual Operations (CIO) from Enacted Intellectual Engagement (EIE), AI-Assisted Performance (AIP), and Demonstrated Independent Capability (DIC), while treating AI Assistance Configuration (AIC) and Situated Opportunity Conditions (SOC) as proximal interpretive conditions rather than learner attributes. Drawing on research on cognitive offloading, human–technology partnership, evaluative judgement, assessment validity, AI literacy, and educational capability, six propositions are formulated for empirical challenge. The contribution is not a new label for offloading, engagement, or independent performance. It lies in linking *ex ante* intellectual responsibility to observed enactment and later demonstrability without assuming that assisted performance and learner capability coincide. This evidentiary layer extends a broader human-centered learning agenda while leaving open a question that cannot be settled in the abstract: which operations should remain attributable to the learner across different disciplines and assessment purposes?

Resumen

La inteligencia artificial generativa ha vuelto menos segura una inferencia educativa habitual: que la calidad de un producto académico constituye evidencia suficiente del trabajo intelectual realizado por el estudiante que lo presenta. Este artículo conceptual desarrolla el Marco de Correspondencia Situada, una arquitectura analítica que conecta tres momentos que a menudo son tratados por separado: las operaciones intelectuales que una actividad asigna al estudiante, la evidencia de que esas operaciones fueron puestas en acto durante una actividad mediada por IA y las capacidades que permanecen demostrables cuando cambian las condiciones de asistencia. El marco distingue la Responsabilidad Epistémica Asignada (AER) y las Operaciones Intelectuales Constitutivas (CIO) del Involucramiento Intelectual Ejercido (EIE), el Desempeño Asistido por IA (AIP) y la Capacidad Independiente Demostrada (DIC), al tiempo que considera la Configuración de la Asistencia de IA (AIC) y las Condiciones Situadas de Oportunidad (SOC) como condiciones interpretativas proximales y no como atributos del estudiante. A partir de investigaciones sobre descarga cognitiva, colaboración entre humanos y tecnología, juicio evaluativo, validez de la evaluación, alfabetización en IA y capacidad educativa, se formulan seis proposiciones susceptibles de contrastación empírica. La contribución no consiste en una nueva denominación para la descarga cognitiva, el

involucramiento o el desempeño independiente. Su aporte reside en vincular la responsabilidad intelectual ex ante con el ejercicio observable y la demostrabilidad posterior, sin asumir que el desempeño asistido y la capacidad del aprendiz coinciden. Esta capa de evidencia amplía una agenda más amplia de aprendizaje centrado en el ser humano, al tiempo que deja abierta una pregunta que no puede resolverse en abstracto: ¿qué operaciones deberían permanecer atribuibles al estudiante en diferentes disciplinas y propósitos de evaluación?

Keywords: generative artificial intelligence; higher education; epistemic responsibility; intellectual engagement; assessment validity; human agency; cognitive offloading

Palabras Claves: inteligencia artificial generativa; educación superior; responsabilidad epistémica; involucramiento intelectual; validez de la evaluación; agencia humana; descarga cognitiva

1. Introduction

Generative artificial intelligence has changed not only what students can produce, but also what a finished product can safely be taken to mean. Explanations, counterarguments, summaries, calculations, conceptual structures, code, evaluations, and polished prose can now emerge within a single interaction. None of this makes AI-assisted work educationally illegitimate. It does, however, loosen a familiar link between the quality of an artifact and the intellectual operations that can be attributed to the student who submits it.

The problem extends beyond academic misconduct. Dawson et al. (2024) argue that assessment validity is more fundamental than cheating because the central educational question is whether available evidence supports the claims institutions make about student capability. Generative AI makes that question more difficult because an observed performance may increasingly belong to a human–AI system rather than to the learner considered independently. Recent validity reasoning further distinguishes problems of construct representation, performance attribution, extrapolation, and assessment use under AI-inclusive conditions (Weidlich, 2026). Coauthorship integrity raises a related question: whether students understand and can defend material generated through human–AI production (Ebrahimzadeh et al., 2026).

Any simple prohibitionist response is complicated by a second body of research. Cognitive offloading is not intrinsically detrimental; humans routinely reorganize intellectual work through external tools (Risko & Gilbert, 2016). Long before generative AI, research on intelligent technologies distinguished improved performance with technology from cognitive effects that remain after interaction with it (Salomon et al., 1991). GenAI sharpens that distinction. Assisted performance may improve without equivalent gains in subsequent no-AI performance, while carefully structured human–AI interaction may sometimes leave learning available after assistance changes (Chen, 2026; Wong & Qiu, 2026). Offloading itself is also heterogeneous: it can reflect dependency, strategic augmentation, or forms of autonomous regulation rather than a single continuum of AI use (Zhu et al., 2026).

The problem is therefore not exhausted by whether students use AI, or by how frequently they use it. A less comfortable question concerns where intellectual responsibility sits during a task, which operations the learner actually performs, and what remains demonstrable once assistance conditions change.

The *Situated Correspondence Framework* is proposed as one way of examining that question. It connects the intellectual operations an educational activity assigns *ex ante* to the learner with evidence of the operations actually enacted during AI-mediated work and with capabilities later demonstrated under specified assistance conditions. Constitutive and auxiliary operations are kept separate; so are the configuration of assistance and the opportunities available to exercise relevant intellectual work. The aim is not to predetermine where legitimate assistance ends. It is to make visible when support, redistribution, and substitution carry different consequences for the educational inference at stake.

The proposal is connected to a broader human-centered learning perspective developed through the *Spectra Human-Centered Learning Framework*, where AI-supported higher education is approached through human agency, active knowledge construction, contextual sensitivity, and socio-technical responsibility (López Ayala, 2026). The present article narrows the lens. Spectra is concerned primarily with conditions of design; situated correspondence asks what evidence is needed before intellectual capability can reasonably be attributed to the learner within such an environment.

That connection is useful precisely because the two problems are not identical. An environment may be designed to preserve agency without revealing which operations a particular learner actually enacted. The reverse is also possible: an observed intellectual operation tells little, by itself, about the pedagogical, technological, institutional, or material conditions that made its exercise possible. Design, enactment, and capability are related, but the framework does not assume that they move together.

The guiding research question is therefore:

How can the correspondence between assigned epistemic responsibility, enacted intellectual engagement, and subsequently demonstrated capability be analyzed in AI-mediated higher education while accounting for the configuration of AI assistance and the situated conditions under which those operations are exercised?

2. Conceptual Design and Evidence-Integration Strategy

This is a conceptual theory-development article. No intervention, participant dataset, experimental effect, or validated psychometric instrument is reported. The methodological task is narrower: to place adjacent theoretical traditions and recent

empirical evidence into a common analytical problem and specify relationships that could later be challenged empirically.

Conceptual research is most useful when theoretical elements are selected because they clarify a defined problem rather than because they exhaust a literature (Jaakkola, 2020). Literature-based theory development likewise requires transparency about the purpose of synthesis and the evidentiary role assigned to different kinds of research (Snyder, 2019). The present study therefore uses theory synthesis and model construction. It is not presented as a systematic review, and it makes no claim of exhaustive database coverage or meta-analytic representativeness.

The synthesis moves across several bodies of scholarship that meet at the attribution problem without fully collapsing into one another. Work on performance, learning, and transfer provides the distinction between what is achieved during activity and what remains available later (Salomon et al., 1991; Soderstrom & Bjork, 2015). Research on cognitive offloading examines how intellectual work is redistributed across human–technology systems (Chen, 2026; Risko & Gilbert, 2016; Zhu et al., 2026). Engagement, self-regulation, AI literacy, and evaluative judgement add learner-level processes (Bearman et al., 2024; Chi & Wylie, 2014; Lan & Zhou, 2025; Long & Magerko, 2020). A separate cluster concerns assessment validity, authorship, judgement, and evidence of learning under GenAI-mediated conditions (Dawson et al., 2024; Ebrahimzadeh et al., 2026; Oliveira et al., 2025; Sharma, 2026; Walton et al., 2026; Weidlich, 2026). Capability and opportunity perspectives complicate the picture further by asking what access can actually be converted into under unequal conditions (Bencsik et al., 2026; Saito, 2026; Walker, 2008).

Recent empirical research from different national settings is used as mechanism-oriented contrast, not as a representative comparison of countries. Its purpose is to examine whether distinctions between supported and retained performance, verification and acceptance, access and capability, or tool fluency and judgement appear under heterogeneous conditions. Country is not treated as a proxy for individual competence, educational quality, or cultural essence.

What follows is accordingly a theoretical proposal, not a validated pathway. The arrows organize claims that can be tested; they should not be read as established causal effects or as a necessary sequence of learner development.

3. Conceptual Foundations: Offloading, Collaboration, and Retained Responsibility

Human cognition has never been tool-free. Writing, diagrams, calculation, databases, search systems, collaborators, and external memory all redistribute cognitive demand. Cognitive offloading names one part of that redistribution: the use of external action or structure to alter the information-processing requirements of a task (Risko & Gilbert, 2016). External support, by itself, is therefore a poor candidate for the educational problem.

The more consequential issue is what the learner no longer needs to perform because support is available—and whether that displaced operation matters to the capability education intends to develop.

This problem precedes generative AI. Salomon et al. (1991) distinguished cognitive effects obtained with a technology from effects of technology that remain after the collaboration. Soderstrom and Bjork (2015) likewise demonstrated why immediate performance can be an unreliable indicator of durable learning. GenAI intensifies that distinction because assistance can now enter operations that previously demanded substantial human interpretation: argument formation, synthesis, explanation, evaluation, restructuring, and reasoning.

Recent evidence reinforces the need to differentiate forms of offloading. Chen (2026) found that performance in AI-assisted writing and subsequent no-AI performance may diverge depending on the cognitive layer transferred to the system. Zhu et al. (2026) distinguish dependent and autonomous cognitive offloading, while Gao and Zhang (2026) identify accountable reconstruction as a mechanism through which learners may recover responsibility for AI-supported material. Wong and Qiu (2026) further show that the temporal ordering of independent thought and AI assistance can matter for subsequent independent outcomes.

This leaves no useful basis for treating offloading as either inherently desirable or inherently suspect. Some operations can be delegated with little consequence for the intended inference. Others cannot. Spell checking may be peripheral to an assessment of causal reasoning; having an AI system generate the causal explanation may not be. Which side of that distinction an operation falls on depends on the construct the activity is meant to evidence.

The analytical focus, then, shifts away from how much AI was used and toward how intellectual responsibility was distributed.

4. From Product to Process: Judgement, Engagement, and Validity

A polished output can be a genuine achievement of a human–AI system and still remain ambiguous evidence of the learner's own capability. Performance quality and educational attribution do not always travel together.

The DRIVE framework shifts attention toward evidence of learning contained in learner–GenAI interactions rather than relying exclusively on the final artifact (Oliveira et al., 2025). Walton et al. (2026) similarly show that university students exercise different forms of judgement while working with GenAI, including judgements concerning knowledge, system limitations, contribution, reliance, and submission. Authentic dialogue data also indicate that engagement with GenAI is heterogeneous and that agentic engagement cannot be inferred merely from the fact that interaction occurred (Luo & Yin, 2026).

Other work illuminates adjacent aspects of the problem. Sharma (2026) foregrounds the visibility of evaluative judgement. Ebrahimzadeh et al. (2026) locate a threat to assessment validity when students submit co-generated material they do not understand. Du and Yuan (2026) conceptualize epistemic dependence in terms of displaced learner judgement and responsibility. Weidlich (2026) asks which inference within the validity chain becomes vulnerable when GenAI changes the conditions under which performance is produced.

These contributions clarify important parts of AI-mediated assessment, yet an unresolved gap remains. Critical interaction with AI does not, by itself, identify which intellectual operations the educational claim required the learner to perform. Nor does understanding a final product necessarily establish that the learner enacted the operations through which the assessed capability was expected to develop.

That gap suggests an earlier point of entry: specify the operations assigned to the learner, look for evidence of their enactment, and only then ask what remains demonstrable.

5. Mechanism-Oriented International Contrast

International research offers no stable basis for a universal story about AI-mediated learning. In Japan, Kawai et al. (2025) reported changes in students' critical-thinking attitudes when ChatGPT-generated counterarguments were incorporated into university learning. The study does not establish generalized gains

in objective critical-thinking ability. Its relevance here is narrower: under some conditions, AI may operate as an intellectual perturbation rather than a substitute for thought.

Research conducted in China likewise reveals heterogeneous mechanisms. Chen (2026) distinguishes AI-supported performance from subsequent no-AI outcomes. Zhu et al. (2026) differentiate autonomous and dependent forms of offloading, while Gao and Zhang (2026) identify verification and reconstruction practices through which doctoral learners make AI-supported work accountable. Lan and Zhou's (2025) systematic review also situates AI use within processes of self-regulated learning rather than equating technological access with learning itself.

Cross-national variation complicates the picture rather than resolving it. Hornberger et al. (2025), studying students in Germany, the United Kingdom, and the United States, identified differences in AI literacy, attitudes, use, and self-efficacy alongside substantial individual variation. National context matters, but the evidence gives little reason to treat it as a cognitive property of the learner.

In Mexico, Velázquez-García et al. (2025) reported performance advantages during ChatGPT-supported educational activity that were not reproduced in the same direction in a subsequent technology-free assessment. The study should not be generalized as a causal law, but it illustrates why assisted achievement and independently demonstrated performance cannot simply be treated as interchangeable evidence.

Other settings contribute different kinds of evidence. In Brazil, de Carvalho et al. (2025) documented frequent use of ChatGPT among 201 medical students alongside concerns about information quality and critical evaluation. Saudi research reports favorable uses and perceptions among medical students, but those perceptions cannot be converted into claims of objective cognitive gain without corresponding performance measures (Almazrou et al., 2025).

In Russia, Masalimova et al. (2026) found critical evaluation to occupy a structurally central position within a psychometric network connecting AI literacy and attitudes. Its relevance is not a claim about Russian learners generally, but evidence that effective engagement with AI contains evaluative dimensions that cannot be reduced to frequency of use.

At a broader level, Bencsik et al. (2026), analyzing 20,242 students across 58 countries, show that national digital infrastructure alone does not determine

functional access to ChatGPT. Institutional, linguistic, technological, and individual conditions intervene. This matters for the present framework because failure to exercise a particular operation cannot be interpreted responsibly without considering whether the learner had a realistic and sustained opportunity to exercise it.

No single international pattern emerges from these studies, and that is analytically useful. Similar levels of AI-supported performance may arise from quite different configurations of access, judgement, reconstruction, offloading, and learner responsibility. The heterogeneity is part of the problem the framework must be able to accommodate, not noise to be averaged away.

6. Conceptual Results: The Architecture of Situated Correspondence


The architecture is deliberately sparse. It separates a central correspondence pathway from an observable system-level outcome and from the conditions needed to interpret both.

Its central theoretical pathway is:

Assigned Epistemic Responsibility / Constitutive Intellectual Operations → Enacted Intellectual Engagement → Demonstrated Independent Capability

A parallel observable outcome is AI-Assisted Performance.

Two proximal conditions—AI Assistance Configuration and Situated Opportunity Conditions—shape how the pathway should be interpreted. A broader institutional, technological, and geopolitical environment functions as boundary context rather than as an individual cognitive construct.

Figure 1*Conceptual Architecture of the Situated Correspondence Framework*

Note. Author's own elaboration. AER = Assigned Epistemic Responsibility; CIO = Constitutive Intellectual Operations; EIE = Enacted Intellectual Engagement; DIC = Demonstrated Independent Capability; AIP = AI-Assisted Performance; AIC = AI Assistance Configuration; SOC = Situated Opportunity Conditions. The relationships represented are theoretical and do not imply established causality.

The figure is not a developmental staircase. Its use is attributional: it keeps visible what the activity intended to assign to the student, what appears to have been enacted, what the human-AI system produced, and what the student can later demonstrate under altered or specified assistance conditions. Those elements may align. They may also diverge.

7. Assigned Epistemic Responsibility and Constitutive Intellectual Operations

Assigned Epistemic Responsibility (AER) designates the intellectual responsibility that an educational task or assessment intends to place on the learner.

Nothing in AER implies a stable psychological trait. It is a property of the educational claim. If an institution intends to claim that a learner can critically

evaluate evidence, responsibility for the evaluative operation must remain meaningfully attributable to that learner.

Constitutive Intellectual Operations (CIO) sit inside that responsibility: they are the operations whose performance by the learner is necessary if the intended educational inference is to remain defensible.

The distinction between constitutive and auxiliary operations is consequential. If the assessed construct is statistical interpretation, formatting a table may be auxiliary while determining what an inferential result means may be constitutive. In another task, however, revision, presentation, or rhetorical judgement might themselves form part of the capability being assessed.

CIO identification must occur before student performance is inspected. Otherwise circularity becomes difficult to avoid: operations visible in successful work could simply be declared constitutive after the fact.

One defensible ordering is:

learning outcome → intended educational claim → operations necessary to sustain that claim → disciplinary expert review → definition of CIO → task design → observed performance

The value of this ex ante specification is methodological rather than rhetorical. It constrains retrospective attribution, although it does not eliminate disciplinary disagreement about what counts as constitutive.

Conceptually, CIO should not be read as a new taxonomy of cognitive skills or as a relabeling of task demands. Existing taxonomies describe kinds or levels of cognition; learning-strategy research describes procedures learners may use; task analysis describes what an activity requires. CIO has a different status. It is an ex ante attributional designation: among all operations involved in a task, it identifies those whose performance must remain attributable to the learner if a particular educational claim is to be warranted. Its novelty therefore lies less in naming intellectual acts than in assigning them a role within an explicit inference from task performance to learner capability.

8. Enacted Intellectual Engagement

Enacted Intellectual Engagement (EIE) concerns evidence that the learner personally performed a task-relevant intellectual operation previously connected to AER and CIO.

That makes EIE narrower than engagement in a general sense. Time on task, interaction length, persistence, emotional involvement, or number of prompts may matter educationally, yet none of them independently demonstrates that a constitutive intellectual operation was performed by the learner.

The distinction also separates EIE from the ICAP framework. ICAP differentiates passive, active, constructive, and interactive modes of engagement through observable learning activity (Chi & Wylie, 2014). EIE asks a different question: whether a particular operation necessary for an intended educational inference can reasonably be attributed to the learner.

EIE also differs from several neighboring constructs. Broad engagement models distinguish behavioral, emotional, and cognitive involvement (Fredricks et al., 2004), while metacognition concerns knowledge about and regulation of cognition (Schraw & Dennison, 1994). Self-regulated learning integrates cognitive, metacognitive, motivational, behavioral, and affective processes used to direct learning (Panadero, 2017). Evaluative judgement, in turn, concerns the capability to judge the quality of one's own and others' work (Tai et al., 2018). Each may support, accompany, or shape intellectual enactment, but none by itself establishes that the specific constitutive operation attached to an educational claim was personally performed.

Accordingly, EIE is neither a global intensity score nor a higher-order learner trait. A highly engaged or self-regulating student could still delegate a particular CIO to AI; a learner could display metacognitive monitoring without executing the operation being monitored; and evaluative judgement would count as EIE only when judgement itself is one of the CIO specified for the task. EIE is therefore indexed to a particular operation, task, and attributional claim. This operation-specificity is the construct's proposed discriminant boundary and an empirical vulnerability: if it cannot be measured separately from these neighboring constructs, the distinction should not be retained.

Evidence of EIE might include reconstructing reasoning, comparing independent sources, rejecting an AI recommendation on defensible grounds, generating an initial interpretation before consulting AI, detecting an error, justifying a decision, or transforming generated material through disciplinary criteria.

Prompt frequency is an especially weak proxy. Sophisticated prompting can orchestrate excellent machine performance while leaving uncertain whether the learner personally enacted the operation under evaluation.

9. AI-Assisted Performance and Demonstrated Independent Capability

AI-Assisted Performance (AIP) names the observable quality of a performance or artifact produced by the human–AI system under specified assistance conditions.

It is treated here as an outcome, not as a latent characteristic of the learner.

That choice blocks one common shortcut. A high-quality AI-assisted artifact cannot automatically be read as high learner capability. The converse is equally important: weak assisted performance may reflect system behavior, access, orchestration, or task–tool mismatch rather than weak underlying capability.

Demonstrated Independent Capability (DIC), by contrast, concerns what the learner can subsequently reconstruct, justify, execute, adapt, or transfer when the available assistance no longer performs the constitutive operation being evaluated.

“Independent” does not mean tool-free cognition. A calculator, dictionary, database, statistical package, assistive technology, or AI system may remain available provided that it does not execute the particular intellectual operation from which the educational inference is being drawn.

Independence is therefore operation-relative, not a claim about cognition in the absence of tools.

Keeping AIP and DIC apart preserves a question that collaborative performance can otherwise obscure: what improved during the interaction, and what intellectual capability remained attributable to the learner afterward?

The EIE–DIC relation should also not be collapsed into the established distinction between performance and learning. Soderstrom and Bjork (2015) distinguish performance observed during acquisition from relatively durable changes in capability reflected in retention or transfer. Situated correspondence accepts that distinction but asks an additional attributional question created by AI mediation: which specific operation embedded in a successful performance was actually enacted by the learner, and can that same operation later be demonstrated when assistance no longer executes it? The unit of analysis is therefore not merely a change from immediate to delayed performance. It is correspondence across an *ex ante* assigned operation, evidence of learner enactment during assisted work, and subsequent operation-relative demonstrability.

Under this formulation, DIC is not synonymous with retention. A learner may retain information without demonstrating the CIO on which an educational claim

depends, and may demonstrate a CIO in a transfer task without reproducing the original product. DIC is tied to the same constitutive operation specified before performance, whereas retention and transfer describe temporal or contextual properties of what is learned more generally. The empirical question is whether evidence of enactment adds explanatory value beyond assisted performance and prior knowledge; P3 is designed to make that distinction testable.

10. AI Assistance Configuration

Quantity of AI use is too coarse to explain how responsibility was redistributed within a task. AI Assistance Configuration (AIC) is introduced as a proximal descriptor of that redistribution.

Four dimensions are provisionally useful: timing, functional scope, generative depth, and constraint or verification structure. Their adequacy remains an empirical question.

Timing concerns whether assistance occurs before a learner's own attempt, during reasoning, after an initial solution, or during revision. Functional scope concerns what AI performs: search, explanation, generation, critique, editing, synthesis, verification, or another operation. Generative depth concerns how deeply assistance enters the intellectual structure of the task, from surface modification to ideational or reasoning-level generation. Constraint and verification structure concerns whether the activity requires independent evidence, justification, source verification, documentation of decisions, or a prior individual attempt.

Two learners can therefore spend comparable amounts of time with AI while inhabiting quite different configurations of intellectual responsibility.

11. Situated Opportunity Conditions

Learning does not occur outside material life. Neither does the opportunity to perform the intellectual work an assessment later attributes to a learner.

Situated Opportunity Conditions (SOC) capture proximal conditions that affect the realistic opportunity, continuity, or cost of performing constitutive intellectual work.

Depending on the research question, relevant conditions may include available study time, employment burden, sleep disruption, connectivity, financial cost,

caregiving responsibilities, linguistic support, technological access, or other factors for which a theoretically plausible mechanism exists.

SOC should not be turned into a generalized adversity index, and constrained circumstances should not become proxies for lower intellectual potential.

The capability perspective is useful here because it distinguishes possession of resources from the substantive opportunity to convert resources into valued forms of functioning (Walker, 2008). Capability-oriented work on generative AI similarly emphasizes that access does not automatically produce meaningful educational agency or equivalent benefit (Saito, 2026).

Its role is interpretive. A learner who had fewer realistic opportunities to exercise a CIO may show lower observed enactment without that pattern warranting a claim about lower underlying potential. The empirical problem is to examine how opportunity, exercise, and later demonstrability relate without treating any one of them as a substitute for the others.

12. Theoretical Propositions and Testable Predictions

The six propositions below make the architecture vulnerable to evidence. They are predictions for empirical challenge, not validated relationships or established causal effects.

P1. Constitutive substitution. Holding the task constant and controlling for prior knowledge, correspondence between AER and EIE will be lower when AIC permits the AI system to execute CIO than when assistance is predominantly restricted to non-constitutive operations.

If that pattern does not appear—if AI execution of CIO leaves learner enactment unchanged—P1 would receive little support.

P2. Verification and defensible judgement. Learners who verify AI-generated claims against independent disciplinary evidence and explicitly justify the criteria used in their evaluation will demonstrate higher-quality defensible judgement than learners whose verification consists primarily of asking generative systems to confirm or reconsider their own outputs, after controlling for prior disciplinary knowledge.

Defensible judgement could be examined through evidentiary relevance, source quality, inferential coherence, explicit warrants, treatment of uncertainty, and expert evaluation. No single indicator is assumed sufficient.

P3. Enactment and subsequent capability. EIE corresponding to CIO will positively predict subsequent DIC after controlling for prior disciplinary knowledge and AIP.

P3 carries the greatest theoretical risk. If constitutive EIE adds no explanatory value to DIC beyond prior knowledge and AI-assisted performance, a central assumption of the architecture would need revision rather than protection by definition.

P4. Configuration beyond frequency. Functional characteristics of AIC—timing, functional scope, generative depth, and verification structure—will explain incremental variance in EIE, AIP, and DIC after controlling for general frequency and duration of AI use.

AIC earns its place only if it adds information. If frequency and duration explain the outcomes just as well, the configuration construct should be reduced or abandoned.

P5. Prompt literacy and capability. Prompt literacy will show a stronger positive association with AIP than with DIC. After controlling for prior disciplinary knowledge and verification quality, its direct association with DIC will diminish substantially.

The empirical issue is whether skill in orchestrating a generative system travels with disciplinary capability or mainly with system-level performance. The framework does not assume the two are unrelated; it predicts that they are not interchangeable.

P6. Situated opportunity and enactment. Greater constraints on available time, continuity of access, or stability of study conditions will be associated with a lower probability of sustained CIO enactment. The relationship between these constraints and DIC will diminish when prior knowledge and EIE are incorporated into the model.

Here the proposed mechanism is opportunity, not fixed intellectual potential.

13. Situated Correspondence: Definition and Unit of Analysis

Situated correspondence is the degree to which evidence indicates that a learner enacted operations previously identified as constitutive of an educational inference and can subsequently demonstrate those operations under specified assistance

conditions, with that trajectory interpreted in light of the opportunities actually available for exercising them.

Two correspondences are of primary interest. AER–EIE asks whether intellectual operations assigned to the student were actually enacted by that student; EIE–DIC asks whether operations evidenced during AI-mediated activity remain demonstrable later.

AIP remains visible in parallel rather than being folded automatically into either correspondence. High AIP can coexist with limited EIE; substantial EIE can accompany modest AIP; both may also be high. Those profiles are theoretically different even when a single performance score would make them look similar.

For that reason, situated correspondence is better treated initially as a profile than as a composite index. Premature aggregation could erase informative combinations of responsibility, enactment, performance, and retained capability.

The choice has a methodological cost: AI use is easy to count, whereas constitutive operations must be identified, their enactment traced, and their later demonstrability tested. The harder measurement problem is also the one on which the framework stands or falls.

14. Methodological Implications

The hardest empirical problem appears before any AI log is coded: identifying CIO with enough reliability to justify later attribution.

For a specified learning outcome, disciplinary experts should articulate the educational claim and independently determine which intellectual operations are necessary to sustain it. Agreement needs to be established before student interaction or performance data are inspected. Otherwise CIO classification risks becoming a retrospective explanation of behavior already observed.

Only then does EIE become an observable coding problem. Potential evidence includes AI interaction records, version histories, oral defenses, stimulated recall, think-aloud protocols, source-verification traces, decision annotations, and structured reconstruction tasks. None is likely to be sufficient across all disciplines.

DIC poses a different measurement problem because the relevant operation must be tested under a condition in which assistance does not perform it. Delayed reconstruction, oral justification, transfer to a novel problem, parallel tasks, or constrained AI configurations are possible designs, not a required battery.

Behavior and interpretation must also be kept apart. A prompt such as “challenge my argument” shows that critique was requested; it does not show that the learner evaluated the critique. Accepting an AI correction is similarly ambiguous unless the grounds for acceptance can be recovered.

Observation can itself disturb what is being observed. Requiring students to document every micro-decision may alter their use of AI, while interaction logs can expose sensitive information, intellectual processes, personal context, or third-party material. Any empirical implementation therefore has to balance evidentiary ambition against reactivity, privacy, consent, data minimization, and institutional governance.

15. Program for Empirical Validation and Framework Testing

There is no single mandatory validation sequence. A practical program can begin where measurement is most uncertain and move only as the preceding evidence allows.

One starting point is operational viability. Disciplinary experts can independently identify CIO for selected tasks; a coding protocol for EIE can then be developed and subjected to inter-rater reliability analysis. Weak agreement at either point would be informative in itself and would require construct revision before more elaborate tests are justified.

A separate empirical route is to manipulate AIC experimentally or quasi-experimentally. Open AI assistance, auxiliary-only assistance, access after an independent attempt, and critique- or verification-only assistance would create contrasting distributions of work and provide direct tests of P1 and P4 under relatively stable task conditions.

Temporal evidence becomes necessary for the central EIE–DIC claim. EIE observed during learning can be related to later DIC while controlling prior knowledge and AIP; SOC can then be examined through longitudinal, within-person, or multilevel designs. These designs need not be confined to one order, but they must preserve the distinction between opportunity, enactment, and later capability.

Cross-national work becomes most informative once the measurement architecture is sufficiently stable to support comparison. Its value would lie in testing invariance, contextual mechanisms, and boundary conditions—not in ranking

national systems or treating country membership as an explanation of individual capability.

16. Discussion

The argument can be stated simply, although its implications are not. Once generative AI can perform substantial portions of intellectual work, the final artifact cannot carry the same attributional burden it once did. That does not invalidate AI-assisted performance. It changes what additional evidence may be needed before capability is assigned to the learner.

Professional cognition is already socio-technical. Engineers use modelling systems, physicians consult decision supports, researchers rely on analytical software, and writers work with increasingly sophisticated digital tools. Higher education has little reason to preserve a fiction of technologically untouched cognition. The difficult boundary lies elsewhere: which operations may legitimately be distributed across a human-AI system, and which must still be demonstrably attributable to the learner for a particular educational claim to remain defensible?

Situated correspondence intersects with work on epistemic dependence, learner-GenAI interaction, coauthorship integrity, evaluative judgement, and validity reasoning (Du & Yuan, 2026; Ebrahimzadeh et al., 2026; Oliveira et al., 2025; Sharma, 2026; Weidlich, 2026). The overlap is substantive, but the point of entry differs. Here the analysis begins with operations assigned *ex ante* to the learner and follows them through enactment and later demonstrability. Dependence, interaction quality, authorship integrity, judgement visibility, and validity each illuminate part of that trajectory without exhausting it.

This is a claim of analytical difference, not superiority. Validity reasoning can identify which inference needs protection; interaction analysis can reveal how a learner works with AI; integrity-oriented approaches can test whether co-generated material is understood; evaluative-judgement research can make decisions visible. Situated correspondence keeps asking a narrower question: did the learner perform the operations required by the educational claim, and can those operations later be demonstrated under specified assistance conditions?

The architecture can also be read alongside the broader human-centered perspective articulated in the *Spectra Human-Centered Learning Framework* (López Ayala, 2026). Spectra foregrounds design conditions under which AI may be

incorporated into learning while preserving human agency, active knowledge construction, contextual sensitivity, and socio-technical responsibility. Situated correspondence does not complete that framework; it opens a different evidentiary problem once learning interactions have occurred.

The two frameworks are connected without being nested in a necessary sequence. Spectra asks about conditions of human-centered AI-mediated learning. Situated correspondence asks what evidence supports attribution of intellectual capability within such learning. Human–AI epistemic behavior cuts across both: learners decide what to trust, verify, delegate, reconstruct, or contest. Those behaviors may help explain correspondence profiles, but no assumption is made that they will do so in the same way across tasks or disciplines.

The same openness is needed for AI orchestration. In some professional and educational contexts, coordinating an AI system may itself be part of the intended capability. Prompt formulation, verification, synthesis, comparison, and delegation decisions could then be constitutive. Under another educational claim, allowing the system to perform those same operations could undermine the inference. Intellectual responsibility is construct-dependent, and the construct is not fixed by the technology.

AIP and DIC may diverge in either direction. A polished human–AI product can coexist with limited later ability to reconstruct the reasoning from which it emerged; a less sophisticated artifact may be followed by stronger capacity to explain, rebuild, adapt, or transfer the relevant operations. Product quality alone cannot resolve that ambiguity.

SOC adds another source of ambiguity rather than a correction factor. Time pressure, unstable access, employment, caregiving, language, economic cost, or technological availability can alter the opportunity to exercise an operation without revealing a learner's underlying intellectual potential. Context matters, but it does not carry a simple deficit meaning.

The educational question therefore shifts from whether AI was used toward where the intellectually consequential work occurred. That question is harder because the answer depends on what a discipline intends to attribute to the learner in the first place. The framework can organize that inquiry; it cannot decide once and for all which operations every field should protect from delegation.

17. Theoretical Contribution

The contribution of the framework is deliberately relational. Cognitive offloading, the distinction between performance and learning, AI literacy, evaluative judgement, epistemic dependence, assessment validity, and the capability approach are not treated as new discoveries. The proposed contribution lies in reorganizing those conversations around a narrower operational problem: the correspondence among intellectual operations assigned to the learner, operations actually enacted during AI-mediated work, and operations that remain demonstrable under specified assistance conditions.

The architecture contributes through separation rather than accumulation. AER and CIO make responsibility explicit before performance is observed; EIE is kept distinct from the human–AI outcome represented by AIP; DIC moves the question to what remains demonstrable later. AIC and SOC qualify interpretation without being converted into learner traits. Whether all of these distinctions survive measurement is left to empirical work.

Taken together, CIO, EIE, and DIC do not claim novelty at the level of cognitive content. Their proposed novelty is structural and inferential: CIO specifies what must be attributable *ex ante*; EIE supplies operation-specific evidence of enactment during AI mediation; DIC tests whether that same operation remains demonstrable under a non-substitutive assistance condition. The three constructs are therefore defined by their relations to an educational claim and to one another, not by asserting that reasoning, engagement, or later performance are newly discovered phenomena.

Within a broader research program on Human-Centered AI in Higher Education, situated correspondence can be read alongside several domains rather than placed at the end of a pipeline. Spectra concerns design conditions for agency and human knowledge construction (López Ayala, 2026). Context-sensitive analysis asks why similar technologies operate differently across institutions and sociocultural environments. Human–AI epistemic behavior examines authority, verification, delegation, and knowledge construction. Situated correspondence focuses on attribution: what was assigned, what was enacted, and what remains demonstrable.

These domains need not unfold in a fixed sequence, nor does one have to be validated before another becomes informative. Their analytical value lies in the possibility of examining how design, context, epistemic behavior, and capability attribution intersect under different empirical conditions.

That openness matters for future research. Particular human-centered designs may or may not produce stronger AER–EIE correspondence; some epistemic behaviors may predict DIC only under certain AIC configurations; SOC may alter opportunities to exercise CIO without behaving as a measure of cognitive potential. The framework becomes useful if such relations can be tested without forcing them into a predetermined order.

Its theoretical value therefore depends less on breadth than on whether attribution can be made more precise, and whether the resulting claims remain vulnerable to disconfirmation.

18. Limitations

The framework remains conceptual. Nothing reported here demonstrates that AER, CIO, EIE, AIP, DIC, AIC, or SOC can yet be measured with sufficient reliability across disciplines, or that the proposed relationships among them will survive empirical testing.

CIO identification is an immediate vulnerability. What counts as constitutive depends on the educational claim, and disciplinary experts may disagree even when that claim appears clear. If operations cannot be specified reliably *ex ante*, one of the framework's central distinctions weakens before any learner data are collected.

EIE raises a different problem: discriminant validity. It must eventually capture something more specific than general engagement, metacognition, evaluative judgement, or interaction quality. If empirical work cannot separate these constructs, terminological distinction would not justify retaining the architecture in its current form.

DIC is no simpler. Any later demonstration is task-dependent, and observed capability may shift with the task, delay, available tools, or assessment mode. Independence must therefore remain relative to the operation and assistance configuration; treating it as an absolute learner property would exceed what the framework can support.

The international evidence also carries unequal weight. Descriptive surveys, qualitative studies, quasi-experiments, conceptual papers, systematic reviews, and multinational samples do not warrant identical inferences. They are used here for theoretical contrast rather than quantitative synthesis. Rapid changes in generative

systems add another instability: the boundary between auxiliary and constitutive support may itself move as systems change.

There is also a methodological tension in making intellectual work visible. Recording interaction or demanding explanation can alter behavior, while AI logs may expose sensitive information, personal context, intellectual processes, or third-party material. More evidence is not automatically better evidence if obtaining it requires disproportionate surveillance.

The connection with Spectra establishes conceptual continuity, not empirical support. A research program becomes more credible only to the extent that its connected components can be measured independently, challenged by competing explanations, revised, or rejected. Continuity should not be mistaken for confirmation.

19. Conclusions

Generative AI has weakened the inference that the quality of an academic product is sufficient evidence of the human intellectual work behind it. That does not make AI-assisted performance intrinsically invalid, nor does collaboration with AI necessarily diminish learning. What has changed is the confidence with which a product can be used, by itself, to attribute capability to a learner.

Situated correspondence separates the pieces of that attribution problem. CIO are identified *ex ante*; EIE asks whether those operations were personally enacted; AIP records the quality of the human–AI performance; DIC examines what remains demonstrable when assistance no longer performs the operation of interest. AIC describes how assistance entered the task, while SOC locates the opportunity conditions under which intellectual exercise was possible. These distinctions may prove useful; they may also prove too fine-grained in some domains.

The framework is therefore better understood as one analytical layer within a wider agenda on human learning mediated by AI, not as the endpoint of that agenda. Design conditions, sociocultural and institutional context, epistemic behavior, and capability attribution are related but not ordered problems. The connection with Spectra lies in their shared concern with human agency, while the present article places greater pressure on the evidence used to say that agency was actually exercised.

Nor can the evaluation problem be reduced universally to “with AI” versus “without AI.” In some settings, competent orchestration of AI may itself constitute legitimate

professional capability. In others, evaluating evidence, constructing causal reasoning, defending an interpretation, or transferring understanding may need to remain personally demonstrable because those operations are the object of the educational claim. The boundary cannot be fixed independently of the construct.

Situated interpretation is equally necessary. Time, work, sleep, connectivity, economic resources, caregiving, language, and institutional support can alter the opportunity and cost associated with intellectual exercise without determining intellectual potential. Context can explain part of a performance trajectory; it should not be converted into a deficit label.

The framework's scientific value depends on whether it can fail. If CIO cannot be identified reliably before performance is observed, if EIE cannot be differentiated from generic engagement, if AIC adds no explanatory information beyond frequency of AI use, or if constitutive EIE does not predict DIC beyond prior knowledge and AIP, the architecture should be revised, simplified, or abandoned.

The question left open for higher education is therefore more precise than how much students use artificial intelligence. It concerns which intellectual operations an institution must be able to attribute to the learner, what evidence shows that those operations were actually exercised, and under what conditions they remain demonstrable as the learner's own capability. Which operations belong in that protected set cannot be decided once for all disciplines. As AI becomes an intellectual intermediary, the distinction among support, distribution, and substitution will have to be argued, evidenced, and periodically reconsidered rather than assumed.

Author Contribution Statement

According to the CRediT Contributor Roles Taxonomy, the authors declare the following contributions:

Mario Guadalupe López Ayala: Conceptualization; Methodology; Formal Analysis; Investigation; Writing – Original Draft; Writing – Review & Editing; Visualization; Project Administration.

Israel Osuna Flores: Conceptualization; Validation; Investigation; Writing – Review & Editing.

Both authors reviewed and approved the final version of the manuscript and assume responsibility for the integrity and academic content of the work.

References

- Almazrou, S., Alanezi, F., Almutairi, S. A., AboAlsamh, H. M., Alsedrah, I. T., Arif, W. M., Alsadhan, A. A., AlSanad, D. S., Alqahtani, N. S., AlShammmary, M. H., Bakhshwain, A. M., Almuhanha, A. F., Almulhem, M., Alnaim, N., Albelali, S., & Attar, R. W. (2025). Enhancing medical students critical thinking skills through ChatGPT: An empirical study with medical students. *Nutrition and Health*, 31(3), 1023–1033. <https://doi.org/10.1177/02601060241273627>
- Bearman, M., Tai, J., Dawson, P., Boud, D., & Ajjawi, R. (2024). Developing evaluative judgement for a time of generative artificial intelligence. *Assessment & Evaluation in Higher Education*, 49(6), 893–905. <https://doi.org/10.1080/02602938.2024.2335321>
- Bencsik, A., Jenei, S., & Módos-Szalai, S. (2026). Who gets to use ChatGPT? A global study on digital access and inequality in higher education. *Social Sciences & Humanities Open*, 13, 102479. <https://doi.org/10.1016/j.ssaho.2026.102479>
- de Carvalho, M. R., Greco, M. E., & de Souza, D. M. (2025). Use of ChatGPT as a study and teaching complementary tool in the medical course. *Educação e Pesquisa*, 51, e288875. <https://doi.org/10.1590/S1678-4634202551288875en>
- Chen, X. (2026). Layer-sensitive cognitive offloading in generative AI-assisted writing: Supported performance and independent no-AI outcomes. *Frontiers in Psychology*, 17, 1906199. <https://doi.org/10.3389/fpsyg.2026.1906199>
- Chi, M. T. H., & Wylie, R. (2014). The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational Psychologist*, 49(4), 219–243. <https://doi.org/10.1080/00461520.2014.965823>
- Dawson, P., Bearman, M., Dollinger, M., & Boud, D. (2024). Validity matters more than cheating. *Assessment & Evaluation in Higher Education*, 49(7), 1005–1016. <https://doi.org/10.1080/02602938.2024.2386662>
- Du, Y., & Yuan, Y. (2026). Epistemic dependence in AI-mediated learning within higher education: A framework for student judgement and responsibility. *Assessment & Evaluation in Higher Education*. Advance online publication. <https://doi.org/10.1080/02602938.2026.2732250>
- Ebrahimzadeh, M., Shibani, A., & Buckingham Shum, S. (2026). Coauthorship integrity: Reconceptualising assessment validity for the age of generative artificial intelligence. *Computers and Education: Artificial Intelligence*, 10, 100609. <https://doi.org/10.1016/j.caeai.2026.100609>
- 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>
- Gao, L., & Zhang, F. (2026). Cognitive offloading and metacognitive calibration in generative AI-mediated doctoral learning: A grounded theory study of accountable reconstruction. *Frontiers in Psychology*, 17, 1920947. <https://doi.org/10.3389/fpsyg.2026.1920947>
- Hornberger, M., Bewersdorff, A., Schiff, D. S., & Nerdel, C. (2025). A multinational assessment of AI literacy among university students in Germany, the UK, and the US. *Computers in Human Behavior: Artificial Humans*, 4, 100132. <https://doi.org/10.1016/j.chbah.2025.100132>
- Jaakkola, E. (2020). Designing conceptual articles: Four approaches. *AMS Review*, 10, 18–26. <https://doi.org/10.1007/s13162-020-00161-0>
- Kawai, R., Ito, T., & Itabashi, S. (2025). Developing critical thinking attitudes in university students using ChatGPT. *Journal of the Japan Society of Technology Education*, 67(1), 51–58. https://doi.org/10.32309/jjste.67.1_51

- Lan, M., & Zhou, X. (2025). A qualitative systematic review on AI empowered self-regulated learning in higher education. *npj Science of Learning*, 10, 21. <https://doi.org/10.1038/s41539-025-00319-0>
- Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376727>
- López Ayala, M. G. (2026). Spectra Human-Centered Learning Framework: A Socio-Technical Architecture for AI-Supported Higher Education. *MenteClara Foundation's Peer-Reviewed Journal*, 11, 441. <https://doi.org/10.32351/rca.v11.441>
- Luo, J., & Yin, X. (2026). Tracing student engagement in dialogues with GenAI in assessment: A semester-long exploratory study. *Innovations in Education and Teaching International*. Advance online publication. <https://doi.org/10.1080/14703297.2026.2663354>
- Masalimova, A. R., Zheltukhina, M. R., Demir, S., Molodozhnikova, N. M., Zvereva, E. V., Berechikidze, I. A., & Lazareva, Y. B. (2026). Mapping human-AI communication in a Russian higher education context: A psychometric network analysis of AI literacy and attitudinal dispositions. *Frontiers in Communication*, 11, 1888712. <https://doi.org/10.3389/fcomm.2026.1888712>
- Oliveira, M., Zednik, C., Bombaerts, G., Sadowski, B., & Conijn, R. (2025). Assessing students' DRIVE: A framework to evaluate learning through interactions with generative AI. *Computers and Education: Artificial Intelligence*, 9, 100497. <https://doi.org/10.1016/j.caeai.2025.100497>
- 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>
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Saito, T. (2026). Redefining agency: A capability-driven research agenda for generative AI in education. *Education Sciences*, 16(1), 155. <https://doi.org/10.3390/educsci16010155>
- Salomon, G., Perkins, D. N., & Globerson, T. (1991). Partners in cognition: Extending human intelligence with intelligent technologies. *Educational Researcher*, 20(3), 2–9. <https://doi.org/10.3102/0013189X020003002>
- Schraw, G., & Dennison, R. S. (1994). Assessing metacognitive awareness. *Contemporary Educational Psychology*, 19(4), 460–475. <https://doi.org/10.1006/ceps.1994.1033>
- Sharma, S. (2026). Educational integrity in GenAI-augmented assessment: Making judgement visible. *International Journal for Educational Integrity*, 22, 6. <https://doi.org/10.1007/s40979-026-00216-y>
- Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104, 333–339. <https://doi.org/10.1016/j.jbusres.2019.07.039>
- Soderstrom, N. C., & Bjork, R. A. (2015). Learning versus performance: An integrative review. *Perspectives on Psychological Science*, 10(2), 176–199. <https://doi.org/10.1177/1745691615569000>
- Tai, J., Ajjawi, R., Boud, D., Dawson, P., & Panadero, E. (2018). Developing evaluative judgement: Enabling students to make decisions about the quality of work. *Higher Education*, 76, 467–481. <https://doi.org/10.1007/s10734-017-0220-3>
- Velázquez-García, L., Cedillo-Hernández, A., Colón-Vallejo, A., & Longar-Blanco, M. del P. (2025). Examining the impact of ChatGPT on Mexican students' performance in higher education. *International Journal of Information and Education Technology*, 15(7), 1368–1376. <https://doi.org/10.18178/ijiet.2025.15.7.2338>

- Walker, M. (2008). A human capabilities framework for evaluating student learning. *Teaching in Higher Education*, 13(4), 477–487. <https://doi.org/10.1080/13562510802169764>
- Walton, J., Bearman, M., Crawford, N., Tai, J., & Boud, D. (2026). How university students work on assessment tasks with generative artificial intelligence: Matters of judgement. *Assessment & Evaluation in Higher Education*, 51(6), 1066–1082. <https://doi.org/10.1080/02602938.2025.2570328>
- Weidlich, J. (2026). Which inference is at risk? Assessment validity reasoning and generative AI. *Assessment & Evaluation in Higher Education*. Advance online publication. <https://doi.org/10.1080/02602938.2026.2734795>
- Wong, S. S. H., & Qiu, S. X. (2026). Think first, ChatGPT later: Guiding human–AI collaboration for learning gains in independent human creativity. *Educational Psychology Review*, 38, 45. <https://doi.org/10.1007/s10648-026-10118-7>
- Zhu, Q., Li, X., Dong, Y., Chang, P., & Fan, M. (2026). Not all cognitive offloading is equal: Distinguishing dependent and autonomous offloading to generative AI. *Frontiers in Psychology*, 17, 1878629. <https://doi.org/10.3389/fpsyg.2026.1878629>