

From Cognitive Augmentation to Learning Alienation: A Conceptual Framework of Generative AI Overdependence in Higher Education

Chen Wang^{1*}

¹ Jiangxi Fenglin College of Foreign Economy and Trade, Jiujiang, Jiangxi, China

Accepted**2026-10-15**

Keywords

GenAI; Learning alienation;
Cognitive delegation;
Overdependence; Epistemic
agency; Higher education

Corresponding Author

Chen Wang

Copyright 2026 by author(s)

This work is licensed under the
CC BY 4.0



<https://doi.org/10.70693/itphss.v3i5.702>

Abstract

A polished assignment reveals little about what its author has learned. This conceptual article examines university students' relationship with academic work when generative artificial intelligence (GenAI) takes on part of it. Through purposive theory synthesis, it connects alienation with distributed cognition, self-regulation, and automation research. Delegating execution differs from delegating evaluation; either can support learning when the activities needed to develop the intended capability remain. Overdependence denotes persistent reliance despite demands or feedback calling for greater participation. The framework distinguishes four forms of estrangement: from reasoning, intellectual direction, developmental meaning, and warranted responsibility for an attributed contribution. Each requires the learner's experience and corroborating evidence, not simply usage counts, dependency scores, or assisted-unaided performance gaps. Nine propositions examine conditional pathways while allowing prior alienation, shared institutional causes, and recovery. The research agenda separates capability from learning relations and compares competing temporal explanations. Teaching proposals center on accessible explanation, revision, and transfer. The framework and its evidence rules remain unvalidated; they are not diagnostic criteria.

1. Introduction

Consider two equally persuasive assignments produced with GenAI. One student checked the evidence and revised the argument; the other submitted an answer they had barely examined. Their work may look similar, but their participation differs. What makes an assisted accomplishment part of learning?

Assistance itself is a poor guide to what students learn. In the college-physics trial reported by Kestin et al. (2025), immediate outcomes favored an instructional AI tutor over classroom active learning. That comparison involved a purpose-built tutor, not unrestricted access to a chatbot.

Fan et al. (2025) separated essay improvement from knowledge and transfer. The ChatGPT group gained more on the essays, but not significantly more on the other two outcomes. Stadler et al. (2024) compared LLMs with search engines: cognitive load was lower with LLMs, yet reasoning and argumentation were weaker. An absent additional benefit is not a relative disadvantage; neither can be inferred from the final product alone.

Over-reliance is reviewed by An et al. (2026) and Zhai et al. (2024). Measurement studies distinguish patterns of reliance (Hou et al., 2025), general dependency (Goh et al., 2025), and programming dependency (Zhang et al., 2026). We retain these distinctions but turn to a question about the completed task: what connects the student to its reasoning, decisions, and outcome?

Alienation concerns an activity that continues despite a sense of detachment (Jaeggi, 2014; Seeman, 1959), a problem also examined in education (Hascher & Hadjar, 2018). In GenAI coursework, we ask whether students understand the reasoning and influence important decisions. What developmental meaning do they find in participating, and can they warrant responsibility for their contribution?

The nine propositions below connect delegation, poorly calibrated reliance, and possible alienation. They also consider earlier estrangement, prior competence, institutional demands, and legitimate divisions of cognitive labor. The standard is meaningful, accountable participation with support—not technological self-sufficiency.

2. Conceptual approach and scope

The central difficulty is to explain detachment without equating external help with lost learning. We use purposive theory synthesis (Jaakkola, 2020) to address it. Distributed cognition and offloading help distinguish productive assistance from substitution within an account of alienation. Agency and self-regulation specify the student's participation; automation research helps assess calibration. Educational GenAI findings bring evidence of benefits and costs to this argument.

The literature was selected for explanatory relevance, not comprehensive coverage. We prioritized 2024–2026 evidence but retained earlier theoretical foundations. Publication details were checked against publisher records, scholarly indexes, and author repositories. Interpretation drew on accessible abstracts or relevant passages; Section 4.2 identifies the preprint. No dual screening, formal quality appraisal, or meta-analysis was undertaken. Findings from other tasks and populations guide the hypotheses but do not validate them.

The unit of analysis is the learner–task–GenAI arrangement across successive learning episodes. The same student may relate differently to work in different courses. Application requires a specified learning objective and an identifiable contribution. The evidence rules describe what to examine, not how an observed relationship arose.

“Unaided” here means without the focal GenAI assistance. Ordinary references and legitimate accommodations remain available; where removal creates a barrier, assessment needs an accessible alternative. Delegation leaves normative responsibilities with learners, educators, and institutions, not the system. Section 9.4 addresses objectives and contributions that are harder to specify.

3. Theoretical foundations

3.1 Alienation as a relation, not a performance deficit

Seeman (1959) distinguishes powerlessness, meaninglessness, normlessness, isolation, and self-estrangement. We adapt rather than reproduce these categories. Powerlessness informs the

concern with intellectual influence. Meaning is separated into understanding the reasoning and valuing participation in learning. A fourth domain concerns responsibility for an attributed contribution. Normlessness and isolation may be antecedents or correlates, but are outside these four task-level domains.

Students can make an idea their own without originating it. Jaeggi (2014) locates appropriation in engagement with, and transformation of, what one encounters, rather than recovery of a fixed authentic self. We extend that distinction to GenAI coursework. The four educational domains are our proposal, not categories supplied by Jaeggi.

Alienation from learning is not a new concern. Hascher and Hadjar (2018) review alienation from learning, teachers, and classmates. Their learning domain overlaps with the process dimension proposed below.

Our focus is narrower: a particular task. A student may value university yet disown a submission, or dislike a course while understanding and directing its work. Existing measures therefore remain important comparators. The proposed domains must explain differences beyond established accounts of alienation, disengagement, or dissatisfaction.

Biesta and Tedder (2007) treat agency as an achievement within an environment. A fixed workflow may leave few intellectual choices; inaccessible support or product-only assessment may narrow them further. These conditions could shape delegation or generate alienation independently of GenAI. Section 7.3 examines both possibilities.

3.2 Distributed cognition and legitimate assistance

Salomon et al. (1991) describe an intellectual partnership with technology in which supported activity has genuine capabilities. Education adds a developmental question: are students acquiring the understanding and influence that the learning purpose requires?

Offloading shifts cognitive demands to external resources (Risko & Gilbert, 2016). In coursework, that shift has to be judged against what the student is meant to learn. Relieving incidental work differs from removing practice in the target capability. Generated material also needs monitoring and evaluation (Tankelevitch et al., 2024), which may leave important work with the learner. Assistance volume does not reveal this allocation.

Effort alone is an inadequate criterion. A beginner may learn by explaining and applying a worked example without first solving the problem unaided. Tetzlaff et al.'s (2025) expertise-reversal meta-analysis supports adjusting assistance to prior knowledge. The educational question is whether that adjustment preserves the learning function.

3.3 Agency, self-regulation, and calibrated reliance

Wu et al. (2025) discuss epistemic agency in GenAI learning partnerships: choosing questions, assessing evidence, and revising conclusions. Generated alternatives may widen students' choices. A record of selections, however, cannot show whether a student weighed the options or simply accepted the system's direction.

Accepting a generated plan does not show how a student assessed it. The record looks much the same whether the learner applied a chosen standard or accepted one supplied by GenAI. Developing evaluative judgment is therefore an assessment concern (Bearman et al., 2024). Guan et al. (2025) review chatbot support for goals, strategies, monitoring, and evaluation, including nongenerative systems. Receiving a plan differs from learning to judge it.

Appropriate trust depends on system capabilities and operating conditions (J. D. Lee & See, 2004). Coursework adds a learning requirement: even accurate assistance can remove needed practice. Educational calibration must address both technical reliability and the capability students are expected to develop.

4. Construct architecture and task–delegation alignment

4.1 Use, reliance, overdependence, and alienation

A use count records interaction with GenAI. Reliance asks whether its assistance entered performance or judgment; dependency asks whether the student feels a persistent need for that help. Goh et al. (2025) and Hou et al. (2025) inform these distinctions in Table 1. A frequent user may still regard assistance as optional, while dependency may be confined to a single operation.

Overdependence names continued reliance on GenAI for learning-central operations when demands or corrective feedback call for fuller participation. The concern is a repeated failure to adjust: suitable attempts are avoided, accessible contradictions go unexamined, or assistance continues unchanged despite the cue. We assess estrangement separately so that its association with reliance can be tested, rather than built into the measure. Neither construct denotes a clinical diagnosis.

4.2 Cognitive delegation and alignment

Who typed the text need not be who made its intellectual decisions. A student can judge a generated draft carefully or type prose that follows the system's reasoning. Execution produces a draft, solution, or interpretation; evaluation sets criteria and decides what to accept. Cognitive delegation allocates this work and its supervision. Although it overlaps with offloading, generated-word counts cannot reveal which judgments students can justify.

Cognitive delegation is not a new term. Di Santi and Florida (2026) use it in a version-4 preprint modeling hybrid performance, reliance, and capability. Their agent-based simulations offer a conceptual comparison, not evidence about university students. Our focus is curricular: the intended capability, regulation of assistance, and appropriation of work.

Grammar correction illustrates why the learning objective matters: it may be peripheral to technical analysis but central to language learning. Extending Biggs's (1996) constructive alignment, task–delegation alignment asks who performs the required work, given the learner's expertise. Substitution occurs when delegation removes a learning-central operation without preserving its developmental function. Reconstructing or applying a generated example can retain that function without an independently produced first answer.

4.3 Performance, capability, and the assistance gap

Salomon et al. (1991) distinguish effects with technology from effects of technology. We apply that distinction to two performance measures: $A(t)$, with the focal GenAI assistance, and $U(t)$, on an equated task without it. Comparable conditions and accommodations give the assistance advantage $G(t) = A(t) - U(t)$. Retention and transfer still need separate tests, as the distinct essay, knowledge, and transfer outcomes in Fan et al. (2025) illustrate.

A large $G(t)$ may accompany effective scaffolding. Unfamiliar tests, time pressure, or difficult items can also depress $U(t)$. Section 8.2 therefore compares independent learning gains across arrangements. Evidence about students' relationship with their work is still needed: a performance gap or learning disadvantage does not establish estrangement.

4.4 Relationship to existing measurement instruments

A scale can describe reliance without explaining a student's part in a particular task. Hou et al. (2025) distinguish reflective, cautious, thoughtless, and collaborative reliance. Goh et al. (2025) measure dependency through cognitive preoccupation, negative consequences, and withdrawal. These scores alone cannot show whether the student understood an inference, influenced its development, or can warrant responsibility for the resulting contribution.

A reliance measure should not contain the alienation outcome it is supposed to predict. Negative-consequence items need scrutiny for such overlap. A study could retain an established

scale's scoring, add calibration observations, and assess alienation separately; scoring adaptations would need validation. Validity tests should establish distinctness and added explanatory value, addressing both discriminant and incremental validity. Which comes first, reliance or estrangement, remains open (Section 5.4).

Table 1: Construct boundaries

Construct	What is assessed	What it cannot establish alone
GenAI use	Academic interaction with GenAI.	Use counts do not establish learning, dependency, or harm.
Reliance	Assistance incorporated into performance or judgment.	Delegation volume does not show calibration.
Dependency	A persistent felt need for, or attachment to, assistance.	Felt necessity is not alienation or a clinical diagnosis.
Overdependence	Repeated reliance on learning-central operations despite demands or corrective cues calling for participation.	One episode, expert efficiency, accommodations, or a dependency score is insufficient.
Learning alienation	Estrangement from reasoning, direction, developmental meaning, or an attributed contribution.	Low performance, reliance, or G(t) needs separate relational evidence (Table 2).

Source: Authors' synthesis of Goh et al. (2025), Hou et al. (2025), Jaeggi (2014), and J. D. Lee and See (2004).

5. Conditional mechanisms and competing trajectories

Propositions P1–P9 set out conditional hypotheses. The following mechanisms allow both sustained augmentation and changes toward, or away from, alienation.

5.1 Augmentation as an enduring possibility

Augmentation can endure when assistance broadens inquiry while preserving learning-central participation. A student might compare two generated explanations, identify the stronger inference, and reconstruct the argument. Neither explanation originated with the student, but evaluating and rebuilding it can retain the learning function. Delegation need not be a stage on the way to dependence.

The ICAP framework distinguishes the quality of cognitive engagement rather than simply counting activity (Chi & Wylie, 2014). A long prompting exchange may involve less constructive thought than a brief consultation followed by critique. Explanation, reconstruction, and application offer more useful observations when matched to the capability being taught.

P1. With comparable task demands and delegated operations, delegation followed by learner evaluation and reconstruction is expected to support stronger retention or transfer than delegation without these activities.

5.2 Substitution and changes in help-seeking

Substitution matters because of the learning opportunity it removes. Delegating an established procedure can free attention for another objective; delegating a procedure still being acquired may remove practice. Novices may nevertheless need scaffolding before independent attempts. P2 concerns a missing learning function, not novice reliance itself.

P2. Substitutional delegation is expected to show a stronger association with later poorly

calibrated reliance under two conditions: the operation is central to the learning objective, and the learner cannot yet adequately evaluate or reconstruct it.

The proposed change concerns the timing of help. After repeated substitution, a student may consult GenAI sooner and make fewer independent attempts. Zhang et al.'s (2026) scale addresses affective, behavioral, and cognitive dependency in programming, but its validation does not demonstrate that sequence. The temporal tests in Section 5.4 must also assess changes in the learning relationship.

5.3 Epistemic deference and evaluative control

An explanation can feel sufficient before its reasons have been checked. Messeri and Crockett (2024) examine illusions of understanding in AI-supported science. Their conceptual analysis concerns a different setting. Applied to coursework, it suggests a risk worth testing: a student may treat the system's persuasive account as a substitute for appraising that account.

Revision is not the same as verification. A student needs independently established criteria and checks of consequential sources to decide whether an answer deserves acceptance. Tankelevitch et al. (2024) discuss the metacognitive work of evaluating GenAI outputs. A revision history shows neither whether checks occurred nor whether they were sound. Delegating production need not remove the learner's evaluative role.

P3. With delegated execution held constant, uncritical delegation of evaluation is expected to predict poorer calibration than learner evaluation using independent criteria and checks of consequential sources.

5.4 Selection, feedback, and reversibility

The direction may also run the other way. A student who already finds an assignment pointless might use GenAI just to finish it. Replacing copied answers with generated ones could leave that relationship unchanged. This hypothetical case represents selection, not technology-induced alienation; Figure 2 gives it equal standing with the forward account.

Baseline marks cannot reveal whether students already felt detached from learning. That relationship needs observation too. Heavy demands, weak preparation, limited support, or output-focused grading might encourage substitution while separately reducing developmental meaning. A same-time correlation cannot distinguish these shared causes from selection, forward effects, or reciprocity. Section 8.3 discusses the design limits.


A possible feedback loop begins when earlier consultation leaves fewer independent attempts. Slower capability growth or altered self-efficacy could then encourage more substitution. An et al.'s (2026) review of dependency traps motivates examining this sequence; it does not establish the temporal links. Observations must distinguish changes in help-seeking from changes in ownership, meaning, and influence.

Help-seeking can change without a loss of meaning or responsibility. Scaffolding might instead build capability and calibrated confidence, reducing substitution and strengthening the learning relationship. Recovery would then be an alternative trajectory in its own right, not alienation whose effects are hidden or delayed.

P4. As substitution recurs, students may seek help earlier and make fewer independent attempts. Changes in capability or self-efficacy could reinforce the pattern. Reciprocal alienation requires both directions within learners: substitution precedes changed alienation, and alienation precedes further substitution. Tests must consider baseline differences and shared causes. Evidence only for the latter direction supports selection; help-seeking changes alone support a behavioral account.


Figure 1 contrasts monitored delegation with recurrent substitution. Figure 2 shows competing temporal explanations; reciprocity requires both directions specified in P4. Autoregressive paths and detailed moderators are omitted for clarity. Instruction, assessment, regulation, and accessible

support may redirect either arrangement.



Instruction, regulation, and assessment can redirect either route.

Figure 1: Conditional learning configurations (Panel A)
Source: Author-developed hypotheses. Context may also affect learning relations without GenAI mediation.



Reciprocal feedback requires both temporal directions (P4).

Figure 2: Competing temporal explanations (Panel B)
Source: Author-developed temporal alternatives; see P4 for the requirement of both directions in reciprocal feedback.

6. Learning alienation: Four proposed dimensions

The four domains concern what students feel estranged from, not four types of student. One, several, or none may apply. Table 2 provides provisional interpretive rules, not validated subscales or diagnostic thresholds. Application requires a specified learning objective, an identifiable contribution, suitable instruction, and accessible opportunities to demonstrate learning.

Each domain requires the student's account and behavioral or contextual support from the same task. Difficulty explaining a result bears on the cognitive question. It does not also establish product alienation, which requires disidentification with the artifact and independent evidence that attribution exceeds warranted responsibility. Record both domains only when each has support. When the evidence remains ambiguous, leave the classification open.

6.1 Cognitive alienation: Estrangement from reasoning

Cognitive alienation concerns reasoning that persistently feels external or opaque despite suitable opportunities to understand it. Reconstruction or adaptation tasks should probe that reasoning directly. Initial unfamiliarity is part of learning; a low score or dislike of an artifact is not enough. Table 2 combines the learner's experience with repeated task evidence.

6.2 Agency alienation: Estrangement from intellectual direction

Understanding a task does not ensure a say in its direction. A student may value the work yet feel unable to affect important goals, methods, or standards. This is the concern of agency alienation. Following Biesta and Tedder (2007) and Wu et al. (2025), we examine actual opportunities to question or change decisions alongside that experience. Accepting GenAI advice after appraisal is different: agreement can itself express agency.

6.3 Process alienation: Estrangement from meaning-making activity

A student may choose and understand a task yet see no value beyond gaining a credential. Process alienation refers to that sustained absence of developmental meaning; it need not begin with GenAI. Fawns's (2022) entangled pedagogy brings purpose, assessment, and technology into the same account. Efficient completion, brief boredom, or limited choice can still coexist with meaningful learning.

6.4 Product alienation: Estrangement from intellectual artifacts

Product alienation concerns identification with, and warranted responsibility for, a specified contribution attributed to the student. It is not a claim about copyright or sole authorship. First establish the student's contribution. Disidentification must then be accompanied by independent evidence that attribution exceeds meaningful decision-making or revision responsibility.

For example, a student might understand a report yet reject responsibility for decisions they could not influence. That differs from finding its reasoning opaque. Disliking the prose, acknowledging assistance, or disputing a conclusion does not by itself establish product alienation.

Different omissions suggest different questions. Lost reasoning practice directs attention to understanding; lost consequential choices direct attention to influence. Rewards confined to output raise a question about developmental meaning, whereas attribution without participation calls for examining ownership. These connections locate the inquiry rather than determine its answer. An omission may affect several relations or may produce no alienation at all.

P5. Recurrent poorly calibrated reliance is expected to relate differently to the four domains, depending on the learning-central operations and participation opportunities it displaces. Tests should account for earlier alienation and context. The four relations require separate assessment from reliance under Table 2; change in one need not appear in the others.

Table 2: Learning alienation: Provisional evidence rules

Domain	Relational object	Provisional evidence requirement	Insufficient evidence / alternatives
Cognitive	Understanding the target reasoning.	Persistent felt opacity or externality, with repeated reconstruction or adaptation difficulty despite suitable learning and assessment opportunities.	One low score, novice unfamiliarity, or dislike. Check instruction, task fit, and access.
Agency	Influence over goals, standards, and intellectual direction.	Felt lack of influence, corroborated by ineffective opportunities to question, override, or reframe consequential decisions.	Agreement, ordinary requirements, or meaninglessness. Consider deliberate acceptance after appraisal.
Process	The developmental meaning of participation.	Sustained absence of developmental meaning, supported by repeated output-seeking or context that detaches participation from learning.	Temporary boredom, efficiency, low effort, or limited choice. Check earlier alienation.
Product	Identification with and warranted responsibility for an attributed contribution.	An identifiable contribution, disidentification with the artifact, and separate evidence that attribution exceeds warranted decision-making or responsibility.	Opacity, generated-word shares, or collaboration. Clarify contributions and revision opportunities.

Source: Author-developed interpretive rules. Corroborate each domain separately; classification does not establish its cause.

7. Boundary conditions and educational design

7.1 Knowledge, literacy, and enacted regulation

AI understanding and evaluation appear in the measures reviewed by Lintner (2024). Subject knowledge is also needed to assess a disciplinary claim. P6 concerns their use in verifying claims, monitoring progress, and adjusting help. These actions carry the proposed protection. Favorable self-ratings and polished outputs cannot replace observations of what students did.

Explicit metacognitive prompts improved aspects of task strategy and self-evaluation in Xu et al. (2025). Their four-week quasi-experiment involved 68 university students, but found no significant achievement difference between groups. Better regulation cannot simply be treated as a general attainment benefit.

P6. Domain knowledge and critical AI literacy are expected to weaken the substitution–overdependence association when learners put them into practice through monitoring, verification, and feedback-responsive adjustment of assistance.

Knowledge, literacy, vigilance, and regulation may describe overlapping checking practices. Studies should identify the actions that matter rather than count favorable scores as separate safeguards. Strong completion incentives may still lead knowledgeable students to substitute

assistance for learning-central activity.

7.2 Assessment and scaffolding

Explaining choices, justifying revisions, and applying ideas to changed problems can make participation visible for feedback (Bearman et al., 2024; Biggs, 1996). Documentation alone may only create more work—and another output that GenAI can produce. Assessment should examine the learning-central judgment behind the record.

P7. Process-sensitive assessment with formative feedback is expected to reduce later alienation relative to product-only assessment when students explain and revise the learning-central elements of their GenAI-mediated work.

Bastani et al. (2025) compared guarded tutoring with unrestricted assistance in high-school mathematics. Tankelevitch et al. (2024) approach GenAI use through metacognitive support. These contributions prompt a design question: what thinking remains with the learner? Difficulty should serve the target capability, not make completion harder for its own sake. Essential scaffolding should remain.

7.3 Institutional incentives, inclusion, and support

Where expectations about GenAI conflict, a student may focus on defending a submission rather than discussing the help used. Consider a course that rewards output over understanding. This hypothetical application draws on Howard et al.'s (2025) meta-analysis of need-supportive and need-thwarting teaching and motivational outcomes; it is not one of its findings. Clear expectations and accessible help could instead support accountable delegation.

When only the final submission receives recognition, critique and reconstruction may become unrewarded work. Clustered deadlines, financial pressures, heavy curricula, or limited teacher access could further constrain participation. These conditions need course-level evidence; they should not be assumed to characterize all higher education.

Course conditions may affect estrangement without delegation. Rigid workflows can limit influence; unclear purposes can weaken meaning. Biesta and Tedder (2007) and Fawns (2022) help make these conditions part of the explanation. GenAI could help students cope or worsen their detachment. Tests must distinguish direct contextual relationships from those mediated by delegation and from conditions that change the strength of a relationship.

P8. Output-centered assessment and perceived performance pressure are expected to strengthen the substitution–overdependence association when independent reasoning receives little recognition and alternative learning support is limited.

Removing GenAI changes more than assistance when the tool also provides access. A student may need language or disability-related support while retaining disciplinary judgment. A lower score after withdrawal could reflect a new barrier, not overdependence. Legitimate accommodations must remain available. When access depends on the focal system itself, an alternative accessible assessment is needed to interpret performance without it.

8. Research agenda: Making the framework vulnerable to evidence

Construct development and baseline design need to proceed together. Measuring reliance and alienation only after a learning arrangement is established cannot separate its effects from earlier selection.

8.1 Establishing discriminant validity

Consistent questionnaire responses do not establish what the items measure. Lintner's (2024) distinction between internal consistency and validity matters here: coherent responses could concern dissatisfaction, incompetence, or dependency rather than alienation. Interviews,

observation, and cognitive interviewing should explore how students interpret candidate items. Include productive collaboration and expanded agency alongside estrangement, testing candidate items against competing interpretations.

Students' accounts of estrangement are necessary, but integrity concerns, defensiveness, or later explanations may shape them. The method presents a further problem when one respondent reports both predictor and outcome in the same format. Their association may partly reflect that common method rather than the relationship under study (Podsakoff et al., 2024).

Ask neutrally about recent episodes under confidential, nonpunitive conditions. Where feasible, separate measurement occasions or sources and code task traces and course conditions independently. A conflict between a student's account and the record deserves investigation; the observer's version should not prevail automatically. Triangulation reveals uncertainties without eliminating bias.

Each item should address one object: reasoning, direction, developmental meaning, or attributed responsibility. Combining them—or naming GenAI as the cause—would assume what the study needs to test. Behavioral evidence should corroborate students' experience, not replace it; reliance observations belong among possible predictors.

Compare correlated domains, higher-order structures, and learner profiles with simpler models of alienation, disengagement, dissatisfaction, or capability. If cognitive and product relations cannot be separated, or agency and process merge, revise the model rather than preserve four factors by assumption.

Establish measurement comparability before comparing groups or time points, then test added explanatory value over existing reliance and dependency measures. For P5, code displaced operations and participation opportunities independently. Otherwise, the coding itself could manufacture the predicted domain pattern.

8.2 Separating assisted performance from capability

Consider a novice who can reconstruct a generated example but cannot yet work fluently alone. A large $G(t)$ may coexist with influence, developmental meaning, and warranted responsibility. This illustrative case separates assistance advantage from alienation; delayed assessment is still needed to establish durable learning.

Let S denote substitution-focused assistance and C augmentation-focused assistance. Compare independent gains as $D = [U(\text{post}) - U(\text{pre})]_S - [U(\text{post}) - U(\text{pre})]_C$. A negative D means that S gained less, even if both groups improved. Random assignment, comparable instructional exposure, delayed transfer tests, and uncertainty estimates strengthen the comparison. Assess any accompanying estrangement through separate relational measures.

P9. With comparable instructional exposure, repeated substitution-focused assistance is expected to produce smaller independent capability gains than augmentation-focused assistance when it removes practice without an equivalent learning activity. This relative disadvantage is expected to be more closely associated with alienation when experienced ownership also declines.

A precise null or favorable result would challenge the prediction. Confidence needs its own assessment: Bastani et al. (2025) found that perceived learning need not match unassisted performance. Capability and learning relations may change differently and need separate evidence.

8.3 Longitudinal and experimental tests

A course can begin after substantial GenAI use. Before the focal arrangement, record prior exposure, alienation, engagement, capability, and motivation, together with task demands, workload, and support. Later recollections can clarify context but cannot replace baseline observations. Confine claims about change to the observed period.

Conventional cross-lagged models may mix between-student differences with within-learner change (Hamaker et al., 2015). Choose observation intervals for changes in practice, capability, and experience. Compare forward-only, reverse-only, reciprocal, and common-cause models. Distinguish baseline confounders from later mediators to avoid adjusting away the proposed pathway. Temporal order and adjustment alone do not establish causation.

An experiment can compare ways of providing help rather than access alone. Conditions might offer hints, examples to reconstruct, unrestricted solutions, or comparable non-GenAI materials. The estimates describe assigned arrangements, not automatically self-selected overdependence. Exposure, explanation quality, implementation, accessibility, and selection need documenting; institutional claims also need course-level evidence. Transfer requires a changed problem, not different wording for the same one.

H.-P. Lee et al. (2025) ask knowledge workers to report critical-thinking practices and effort; this is not experimental evidence of lost student capability. Bastani et al. (2025) experimentally compare tutoring arrangements, but in high-school mathematics. Neither measures university learning alienation. These neighboring studies inform designs without settling the present questions.

8.4 Competing explanations and disconfirmation

Prior alienation, preparation, motivation, course conditions, and support are substantive alternatives. An apparent cost may also reflect assessment mismatch or useful redistribution toward another capability. Table 4 turns these possibilities into design priorities. A reverse-only alienation-to-substitution pathway supports selection rather than reciprocity.

Revise the framework if substitution repeatedly supports durable capability and ownership, the domains add nothing to established constructs, or well-implemented interventions leave predicted relationships unchanged. Interpret these results with adequate precision and implementation evidence; nonsignificance alone is inconclusive. Table 3 specifies a distinct test for each proposition.

Table 3: Nine propositions and their tests

ID	Proposed relationship	Required contrast	Relevant observations
P1	Evaluation and reconstruction improve retention or transfer after delegation.	Comparable operations, demands, and assistance quality.	Critique, reconstruction, delayed explanation, and transfer.
P2	Task centrality and limited evaluative capability strengthen the substitution–overdependence link.	Missing learning activity versus appropriate novice scaffolding.	Learning targets, prior capability, and responses to corrective cues.
P3	Uncritical evaluative delegation predicts poorer calibration beyond execution delegation.	Separate production from judgment.	Independent criteria, source checks, and reasoned correction.
P4	Behavioral reinforcement may sustain substitution; reciprocal alienation requires both directions.	Behavioral-only, forward, reverse, reciprocal, and shared-cause accounts.	Consultation timing, attempts, lagged relational measures, capability, and self-efficacy.

ID	Proposed relationship	Required contrast	Relevant observations
P5	Relational change varies with displaced operations and participation.	Code displacement independently; assess each domain separately.	Domain-specific accounts, reasoning, choices, activity, and contribution records.
P6	Knowledge and literacy protect through enacted regulation.	Observed monitoring versus self-ratings alone.	Verification and feedback-responsive adjustment.
P7	Process-sensitive assessment and feedback reduce later alienation.	Product-only assessment; comparable implementation.	Relational change, explanation, revision, access, and burden.
P8	Pressure and output incentives strengthen the substitution–overdependence link.	Variation in recognition of reasoning and other support.	Expectations, reward criteria, concentrated demands, and calibration.
P9	Substitution without equivalent learning activity produces smaller independent gains.	Augmentation-focused assistance; assess relational outcomes separately.	Equated baseline, post-task and delayed capability, transfer, and ownership.

Source: Author-developed hypotheses and tests. P4 requires both temporal directions; P5 requires separate evidence for each domain.

Table 4: Research priorities and competing evidence

Priority	Design requirement	Evidence challenging the proposal
Baseline / selection	Observe earlier alienation, use, engagement, capability, motivation, demands, and support. Separate baseline causes from later mediators.	Prior alienation or shared conditions explain the association without an additional forward effect.
Construct validity	Combine interviews, independent relational and behavioral evidence, reporting-bias precautions, and established measures.	No added explanation, or associations depend on shared reporting methods.
Calibration	Track execution, evaluation, consultation timing, attempts, and feedback-responsive adjustment across comparable tasks.	Purpose and expertise explain differences without a remaining substitution relationship.
Learning contrast	Vary assistance organization, equate exposure, and separately assess baseline and delayed capability, transfer, and relations.	Substitution repeatedly supports lasting capability and meaningful appropriation.
Temporal direction	Compare forward, reverse, reciprocal, and shared-cause models; separate within-person change from stable differences.	Only prior alienation predicts later substitution.

Priority	Design requirement	Evidence challenging the proposal
Institutional context	Observe expectations, grading, workload, and support; distinguish direct, mediated, and moderating relationships.	Context explains alienation without delegation, or predicted incentives show no association despite adequate variation.
Intervention	Compare process-sensitive feedback with product-only assessment; assess implementation, access, burden, and selection.	A well-implemented intervention leaves the predicted relation unchanged with adequate precision.
Scope / inclusion	Specify objectives and contributions, preserve accommodations, and test measurement comparability.	Access barriers, unclear attribution, or incomparable tasks explain the apparent alienation.

Source: Author-developed research priorities. Interpret null results against precision and implementation, and limit claims to observed objectives and exposure.

9. Discussion and implications

9.1 What the framework adds

The framework distinguishes understanding, influence, developmental meaning, and warranted responsibility at the task level, separately from reliance. Section 4.4 and Table 2 make that distinction testable. Without added explanatory value over established measures, the framework would merely rename familiar concerns.

Where earlier alienation or shared course conditions explain substitution, restricting GenAI leaves the underlying problem in place. The educational response would need to address purposes, meaningful choices, workload, and feedback (Biesta & Tedder, 2007; Fawns, 2022). This account does not exclude a further effect of substitution: selection and forward effects may operate together.

9.2 The augmentation–alienation paradox, qualified

GenAI can explain an operation while making it unnecessary to perform that operation to submit work. Vaccaro et al. (2024) find task-dependent benefits and limitations in human–AI combinations. Their meta-analysis concerns task performance, not durable educational capability or ownership. Better output therefore does not resolve this educational tension.

Better explanations and feedback may accompany easier substitution for students' work. Section 1's tutoring evidence keeps the beneficial possibility in view. Fawns (2022) places educational effects in the interaction of purposes, pedagogy, and technology. Whether stronger systems support learning depends on how assistance is organized, rather than a fixed property of GenAI.

9.3 From detection to accountable participation

Teaching should make students' contributions visible through reasoned choices, purposeful revisions, and transfer. These demands need to be accessible, proportionate, and useful for feedback. Their value lies in the thinking they reveal, not in how difficult they are to automate. Otherwise, documenting a process becomes another compliance task.

Delegation literacy is an instructional aim: students identify what they must learn, allocate work, check consequential claims, and evaluate learning. We draw on AI literacy and

self-regulation (Guan et al., 2025; Lintner, 2024), without claiming a validated additional competence. Two possible teaching designs illustrate these judgments.

An argumentation exercise could start with a student's claim and a generated response. The student might question an unsupported premise and explain its implications for the conclusion. A different case would probe transfer; an oral explanation or accessible equivalent would examine a central judgment. Assessment would follow the reasoning through these changes, not a record of completed steps.

In programming, students could inspect generated code, predict its behavior, find an input on which it fails, and justify a revision. For example, what happens when the input list is empty? A novice may first need a worked example. Changing a requirement and asking students to defend a design choice goes beyond checking whether the program runs. Responsibility concerns the assigned logic, not how much code was typed manually.

Students cannot meaningfully justify a choice that cannot affect their work. Such choices need clear expectations and recognition of reasoning, alongside coordinated workloads and accessible alternatives. Evaluation should ask what course organization allowed students to do as well as how they experienced participation. More responsibility alone does not create more influence.

9.4 Limitations and boundary cases

Relevant perspectives may have been missed in this selective synthesis. Its four domains may overlap in practice, and short-term findings cannot settle longer trajectories. Both the framework and teaching responses await validation. Testing may justify merging domains, revising evidence rules, or restricting scope.

Application also requires a defensible learning objective and an identifiable contribution. Language polishing or access support is not automatically substitution. Three illustrative boundary cases show why these limits matter.

Graduate original research. An evolving doctoral project may be too broad for this framework, whereas defending an analytical method supplies a bounded training task. Delegating routine coding leaves method-selection responsibility with the researcher. Persistent opacity after suitable instruction warrants examining the cognitive relation. Product alienation still needs separate evidence under Table 2; uncertainty in exploratory work alone establishes neither.

Group projects. Interpretation begins by locating responsibility. A literature reviewer need not reproduce a teammate's code. If an unreviewed GenAI synthesis is credited to that reviewer, examine identification with it, relevant decisions, and opportunities to revise. Understanding the reasoning remains a separate cognitive question. Neither disagreement nor division of labor is sufficient evidence; unclear attribution leaves the case unresolved.

Artistic creation. Students may intentionally seek surprise, opacity, or distributed authorship. With a specified goal such as compositional judgment or critical reflection, curating generated material and explaining aesthetic choices may be the relevant contribution. Reproducing the model's internal process is a different requirement. Exploration without an identifiable learning target falls outside the framework; sole authorship and technical self-sufficiency are not universal standards.

Factual errors, plagiarism, and privacy breaches may accompany alienation but do not establish it. Accurate, compliant work may still feel estranged. Nor can empirical findings alone decide which capabilities universities ought to require. That remains an educational judgment, beyond demonstrating that a task can be completed without technological help.

10. Conclusions and suggestions

The student's continuing part in the work is the focus of this framework. Understanding, intellectual influence, developmental meaning, and warranted responsibility are connected to delegation and calibration through nine propositions. Tests must distinguish these relations from performance and dependency, allowing selection, shared causes, and recovery. For teaching, the proposal is to leave accessible opportunities to explain, revise, and transfer ideas. These sustain the aim of accountable participation; reducing GenAI use is not itself the objective.

References

- An, C., Lu, K., & Yang, J. (2026). The dependency trap: A systematic review of the contributing factors and consequences of student over-reliance on generative AI in higher education. *Humanities and Social Sciences Communications*, 13, Article 1604. <https://doi.org/10.1057/s41599-026-08959-2>
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), Article e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- 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>
- Biesta, G., & Tedder, M. (2007). Agency and learning in the lifecourse: Towards an ecological perspective. *Studies in the Education of Adults*, 39(2), 132–149. <https://doi.org/10.1080/02660830.2007.11661545>
- Biggs, J. (1996). Enhancing teaching through constructive alignment. *Higher Education*, 32(3), 347–364. <https://doi.org/10.1007/BF00138871>
- 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>
- Di Santi, E., & Florida, C. (2026). *Cognitive amplification vs cognitive delegation in human-AI systems: A metric framework* (Version 4) [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2603.18677> Version 4: <https://arxiv.org/abs/2603.18677v4>
- Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, 56(2), 489–530. <https://doi.org/10.1111/bjet.13544>
- Fawns, T. (2022). An entangled pedagogy: Looking beyond the pedagogy–technology dichotomy. *Postdigital Science and Education*, 4(3), 711–728. <https://doi.org/10.1007/s42438-022-00302-7>
- Goh, A. Y. H., Hartanto, A., & Majeed, N. M. (2025). Generative artificial intelligence dependency: Scale development, validation, and its motivational, behavioral, and psychological correlates. *Computers in Human Behavior Reports*, 20, Article 100845. <https://doi.org/10.1016/j.chbr.2025.100845>
- Guan, R., Raković, M., Chen, G., & Gašević, D. (2025). How educational chatbots support self-regulated learning? A systematic review of the literature. *Education and Information Technologies*, 30(4), 4493–4518. <https://doi.org/10.1007/s10639-024-12881-y>
- Hamaker, E. L., Kuiper, R. M., & Grasman, R. P. P. P. (2015). A critique of the cross-lagged panel model. *Psychological Methods*, 20(1), 102–116. <https://doi.org/10.1037/a0038889>

- Hascher, T., & Hadjar, A. (2018). School alienation—Theoretical approaches and educational research. *Educational Research*, 60(2), 171–188. <https://doi.org/10.1080/00131881.2018.1443021>
- Hou, C., Zhu, G., Sudarshan, V., Lim, F. S., & Ong, Y. S. (2025). Measuring undergraduate students' reliance on generative AI during problem-solving: Scale development and validation. *Computers & Education*, 234, Article 105329. <https://doi.org/10.1016/j.compedu.2025.105329>
- Howard, J. L., Slemp, G. R., & Wang, X. (2025). Need support and need thwarting: A meta-analysis of autonomy, competence, and relatedness supportive and thwarting behaviors in student populations. *Personality and Social Psychology Bulletin*, 51(9), 1552–1573. <https://doi.org/10.1177/01461672231225364>
- Jaakkola, E. (2020). Designing conceptual articles: Four approaches. *AMS Review*, 10(1), 18–26. <https://doi.org/10.1007/s13162-020-00161-0>
- Jaeggi, R. (2014). *Alienation* (F. Neuhouser, Ed.; F. Neuhouser & A. E. Smith, Trans.). Columbia University Press. <https://doi.org/10.7312/jaeg15198>
- Kestin, G., Miller, K., Klales, A., Milbourne, T., & Ponti, G. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports*, 15, Article 17458. <https://doi.org/10.1038/s41598-025-97652-6>
- Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (Article 1121, pp. 1–22). Association for Computing Machinery. <https://doi.org/10.1145/3706598.3713778>
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
- Lintner, T. (2024). A systematic review of AI literacy scales. *npj Science of Learning*, 9, Article 50. <https://doi.org/10.1038/s41539-024-00264-4>
- Messeri, L., & Crockett, M. J. (2024). Artificial intelligence and illusions of understanding in scientific research. *Nature*, 627, 49–58. <https://doi.org/10.1038/s41586-024-07146-0>
- Podsakoff, P. M., Podsakoff, N. P., Williams, L. J., Huang, C., & Yang, J. (2024). Common method bias: It's bad, it's complex, it's widespread, and it's not easy to fix. *Annual Review of Organizational Psychology and Organizational Behavior*, 11, 17–61. <https://doi.org/10.1146/annurev-orgpsych-110721-040030>
- 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>
- 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>
- Seeman, M. (1959). On the meaning of alienation. *American Sociological Review*, 24(6), 783–791. <https://doi.org/10.2307/2088565>
- Stadler, M., Bannert, M., & Sailer, M. (2024). Cognitive ease at a cost: LLMs reduce mental effort but compromise depth in student scientific inquiry. *Computers in Human Behavior*, 160, Article 108386. <https://doi.org/10.1016/j.chb.2024.108386>
- Tankelevitch, L., Kewenig, V., Simkute, A., Scott, A. E., Sarkar, A., Sellen, A., & Rintel, S. (2024). The metacognitive demands and opportunities of generative AI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Article 680, pp. 1–24). Association for Computing Machinery. <https://doi.org/10.1145/3613904.3642902>
- Tetzlaff, L., Simonsmeier, B., Peters, T., & Brod, G. (2025). A cornerstone of adaptivity – A

meta-analysis of the expertise reversal effect. *Learning and Instruction*, 98, Article 102142. <https://doi.org/10.1016/j.learninstruc.2025.102142>

Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12), 2293–2303. <https://doi.org/10.1038/s41562-024-02024-1>

Wu, J.-Y., Lee, Y.-H., Chai, C. S., & Tsai, C.-C. (2025). Strengthening human epistemic agency in the symbiotic learning partnership with generative artificial intelligence. *Educational Researcher*, 54(6), 358–368. <https://doi.org/10.3102/0013189X251333628>

Xu, X., Qiao, L., Cheng, N., Liu, H., & Zhao, W. (2025). Enhancing self-regulated learning and learning experience in generative AI environments: The critical role of metacognitive support. *British Journal of Educational Technology*, 56(5), 1842–1863. <https://doi.org/10.1111/bjet.13599>

Zhai, C., Wibowo, S., & Li, L. D. (2024). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: A systematic review. *Smart Learning Environments*, 11, Article 28. <https://doi.org/10.1186/s40561-024-00316-7>

Zhang, H., Yang, Y., Zhang, Y., Qiu, B., & Yang, J. (2026). Generative AI dependency on programming among university students: A scale development and validation study. *BMC Psychology*, 14, Article 1245. <https://doi.org/10.1186/s40359-026-05011-5>