ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

What AI can and cannot do: A cognitive-depth framework for differentiated AI integration in higher education

Michael Lin

The University of Queensland

Higher education has responded to generative AI largely by regulating it by task, specifying where the technology may or may not be used in each assessment. This paper argues that task-based regulation uses the wrong unit of analysis. A single task typically spans several types of cognitive work, and a tool that can be welcomed for one type may quietly displace another. This conceptual paper proposes a four-level framework of cognitive engagement, adapted from classical Chinese expressions about knowing and articulated in contemporary cognitive terms: broad knowledge acquisition, analogical transfer, contextual inference, and original creation. The levels are ordered along a single axis, from cognitive work that is codifiable and reducible to rules and patterns, to work that is situated and dependent on human judgement. This axis maps closely onto the boundary of what generative AI can and cannot do: capability is near-complete at the codifiable end and degrades sharply as work becomes situated. The framework thus functions as a diagnostic lens for design, allowing educators to assign AI a role matched to the cognitive work a task is meant to develop. Two illustrative cases show the framework guiding differentiated, defensible decisions about AI integration within a single assessment.

Keywords: artificial intelligence, cognitive frameworks, assessment design, academic integrity, generative AI, conceptual analysis

1. Introduction

Higher education responses to generative AI tend to fall into two broad positions (Chan, 2023; Xiao et al., 2023). The permissive position treats the technology as an opportunity to be embraced, encouraging students to use AI across their learning in the belief that these tools are now part of professional practice and that withholding them leaves graduates unprepared (Chan, 2023). The restrictive position treats it as a risk to be contained. AI use is restricted or prohibited in the name of academic integrity, assessment security, and the assurance demands of external accreditation. These concerns have prompted national regulators to call for assessment reform (Lodge et al., 2023). Each position has an internal logic and serious advocates.

Most institutions have adopted a graduated middle rather than occupying either pole in its pure form, issuing increasingly detailed guidance that permits AI for some tasks and forbids it for others, refined assessment by assessment (Chan, 2023; Xiao et al., 2023). Empirical analysis of institutional policies bears out this pattern (Xiao et al., 2023). The graduated middle, however, differentiates along the wrong dimension. It sorts tasks by whether AI is permitted, when what determines whether AI use is sound is not the task but the kind of cognitive work the task requires. Even where such guidance is detailed, drawing a defensible line around acceptable AI use has proved persistently difficult in practice (Corbin et al., 2025). An analytical report, for instance, may ask students to gather and summarise material, apply a method to a new situation, interpret incomplete evidence, and construct an original argument. A task-level rule cannot differentiate among these forms of cognitive work, however detailed it is; its categories are at the wrong granularity.

Defining AI use by task also means that task-level rules must continually expand and are rarely up to date as new tasks emerge. Regulating a pattern-based technology with an ever-expanding set of task-based patterns is a race that cannot be won. The more useful question is therefore not whether students should use AI on a given task, but a prior one that the debate skips over: what kinds of cognitive work can a learner soundly delegate to AI, and what work cannot be delegated without the learner losing something essential to the capability the task was meant to build? Neither embrace, prohibition, nor the graduated middle answers this question, because all three lack a way of distinguishing one kind of cognitive work from another. Educators meet this gap whenever they design an assessment and decide what a student may reasonably delegate to an

AI system and what they must do unaided in order to learn. Learning theory provides substantial insight into how students learn, but it has not produced a widely adopted framework for locating the boundary between AI support and AI substitution. Without one, the decision falls back on individual intuition, applied unevenly from task to task, at precisely the moment when observable, genuine learning has become one of the defining questions of the field.

This paper argues that the line can be drawn in a principled way, starting from an observation the task-level debate obscures: generative AI's capability is profoundly uneven across kinds of cognitive work. On some it already performs at or beyond human level; on others it fails, and fails not for want of refinement but structurally, in ways that further training is unlikely to close. Any approach that treats this steeply graded capability as though it were uniform, whether by embracing AI, prohibiting it, or rationing it by task, is calibrated to the wrong variable. What the unevenness calls for is differentiation by cognitive work: an integration strategy that asks of each piece of learning what kind of thinking it requires, and matches the role of AI to that.

To make that differentiation possible, this paper proposes a four-level framework of cognitive engagement, adapted from classical Chinese expressions about knowing but recast in contemporary terms. The framework orders cognitive work along a single axis, from the codifiable, reducible to rules and patterns, to the situated, which depends on judgement that resists such reduction. This axis maps closely onto the boundary of what generative AI can and cannot do, and this correspondence lets the framework serve as a practical lens: a way of seeing, for any piece of learning, whether AI can soundly support it or would displace cognitive work the learner needs to develop.

The paper is conceptual and its method is conceptual analysis. The four levels are derived from the classical expressions and specified in contemporary cognitive terms (Section 3), mapped against the literature on generative AI capability (Section 4), and translated into design guidance illustrated through two cases from undergraduate business education (Section 5). The work involved no data collection and no human participants, and ethics approval was therefore not required. Section 6 sets out the limitations that follow from this scope.

2. Background and related work

2.1 Approaches to AI integration in higher education

Institutional and scholarly responses to generative AI have grown rapidly (Crompton & Burke, 2023), reshaping assessment across the sector (Xia et al., 2024). They can be grouped into three broad forms. The first is tool-centred: guidance on which AI tools are sanctioned, how to access them, and how to prompt them effectively, on the premise that capability with the tools is itself the graduate outcome. The second is policy-centred: statements of permitted and prohibited use, oriented largely toward academic integrity and assessment security. The third, and increasingly mainstream form, is the graduated, task-by-task guidance described in Section 1, which specifies AI's permissibility assessment by assessment, often within an institution-wide policy framework (Chan, 2023).

What these approaches share is an orientation toward the task or the tool as the unit of analysis. Few take the kind of cognitive work that a task elicits as their primary unit of analysis, and so they are not positioned to address the difficulty identified in Section 1: that a single task typically spans several kinds of cognitive work with different implications for AI use.

2.2 Frameworks of cognition and learning

The obvious place to seek a cognitive unit of analysis is the established literature on learning frameworks. The most influential is Bloom's taxonomy (Bloom et al., 1956) and its later revision (Anderson & Krathwohl, 2001), which orders educational objectives from remembering and understanding through to evaluating and creating. Related schemes, such as the SOLO taxonomy (Biggs & Collis, 1982), similarly describe ascending levels of structural complexity in student understanding.

These frameworks have been enormously productive, but for a particular purpose. They were designed to specify and assess learning outcomes, breaking cognition into discrete, orderable, measurable categories aligned to curriculum and assessment. That strength is also why they do not answer the present question: outcomes are operationalised into observable, standardised performances judged against explicit criteria, the logic of constructive alignment that has shaped outcome-based assessment (Biggs, 1996). A performance specified clearly enough to be marked reliably is, by that fact, explicit and rule-like, and so codifiable, and a pattern-based system can increasingly simulate it. A framework whose levels are defined by their measurability is thus oriented toward observable outputs, and observable outputs are precisely what generative AI can increasingly supply without performing the cognitive work behind them. Such a framework is poorly placed to mark the boundary of what AI lacks: the point at which cognition stops being cleanly decomposable and must be grasped as a situated whole. Recent work has begun to revise Bloom's taxonomy for an AI context (Gonsalves, 2026); the present paper takes a different route, holding that the limitation cannot be revised away because it follows from the taxonomy's analytic, measurement-oriented form.

2.3 An analytic lens and a holistic one

This marks a difference in kind. Bloom and SOLO are analytic: they decompose cognition into component levels that can be specified and tested. The framework introduced next is holistic: each level is a gestalt, an integrated mode of cognitive work captured in a compact image rather than a decomposed, measurable category. This contrast is itself culturally rooted; comparative research associates holistic cognition with Chinese intellectual traditions and analytic cognition with Western ones (Nisbett et al., 2001). Neither is superior; they serve different ends. For the present analysis, correct output is not treated as sufficient evidence of understanding, and a boundary of this kind cannot be located by decomposing cognition into measurable parts because the decomposition captures precisely the component that can be reproduced without understanding.

2.4 The stakes: AI use and the formation of capability

A final strand of work shows why the boundary matters. Research on cognitive offloading (Risko & Gilbert, 2016) predates generative AI but gains new urgency with it. Its classic finding is well established: when cognitive work is delegated in a way that replaces rather than extends a person's own processing, the capability that work would have exercised tends not to develop or consolidate. How far this extends to generative AI specifically is still being established, but early studies raise the concern that habitual reliance on AI may weaken the opportunities through which independent judgement is formed. A systematic review links over-reliance on AI dialogue systems to declines in students' critical thinking and analytical reasoning (Zhai et al., 2024), and in one mixed-methods study, frequent AI-tool use was associated with weaker critical thinking, mediated by increased cognitive offloading (Gerlich, 2025).

2.5 The gap

Together, these literatures leave a specific gap. Approaches in AI integration govern tools and tasks rather than cognitive work; established learning frameworks organise cognitive work for measurement rather than for locating the boundary of machine capability; and cognitive-offloading research establishes the stakes without identifying that boundary. What is missing is a framework that distinguishes where AI can support learning from where it would displace the cognition education is meant to develop.

3. The four-level cognitive framework

3.1 Overview

The four levels, Broad Knowledge Acquisition and Retention (博闻强识), Analogical Transfer (触类旁通), Contextual Inference (一叶知秋), and Original Creation (无中生有), describe an ascending progression in the depth of cognitive work, from the accumulation of knowledge to the generation of genuinely novel ideas. All four are four-character idioms (*chengyu*) still in everyday use in modern Chinese, each compressing a mode of knowing into a single image; they belong to a classical tradition whose distinctive cultural model of learning is

well documented (Li, 2012), and the framework draws on the images rather than on their philology. Table 1 summarises the four levels and previews the relationship between each and the capabilities of generative AI developed in Section 4. Section 6 qualifies the final column.

Table 1.
The four-level cognitive framework and its relationship to current and near-term generative AI capability

Level	Classical term	Literal image	Cognitive term	Defining characteristic	AI capability (current and near-term)
1	博闻强识 (bó wén qiáng zhì)	broad learning, strong retention	Broad Knowledge Acquisition and Retention	Accumulation and faithful reproduction	Near-complete coverage
2	触类旁通 (chù lèi páng tōng)	touching one category, comprehending others	Analogical Transfer	Cross-contextual mapping	Substantial but unreliable
3	一叶知秋 (yī yè zhī qiū)	one leaf, knowing autumn has come	Contextual Inference	Judgement under incompleteness	Partial at best; degrades sharply
4	无中生有 (wú zhōng shēng yǒu)	producing something from nothing	Original Creation	Genuine novelty	Beyond current and near-term capability

3.2 Level 1: Broad knowledge acquisition and retention (博闻强识)

The classical term 博闻强识 (bó wén qiáng zhì) means broad learning and strong retention. It names the foundational cognitive work of acquiring information widely and retaining it reliably for recall. The learner accumulates facts, concepts, terminology, and procedures, and reproduces them accurately on demand. The demands are real but bounded: success depends on breadth and accuracy rather than situated judgement. Higher-level work depends on this retained knowledge, but at Level 1 knowledge need only be correct and available, not transformed or applied.

3.3 Level 2: Analogical transfer (触类旁通)

The classical term 触类旁通 (chù lèi páng tōng) means, literally, touching one category and comprehending others. It names the capacity to take understanding gained in one context and apply it productively in another. The learner recognises structural similarities across superficially different situations and transfers a principle, method, or insight from a familiar domain to an unfamiliar one. The defining characteristic of this level is cross-contextual mapping. In cognitive terms, this is analogical reasoning grounded in relational structure rather than surface similarity (Gentner, 1983). Crucially, the quality of analogical transfer depends on grasping why a principle holds, not merely that it was used before. Transfer grounded only in surface similarity, without causal understanding, produces superficially plausible but unreliable results, a distinction that becomes important when we turn to AI.

3.4 Level 3: Contextual inference (一叶知秋)

The classical term 一叶知秋 (yī yè zhī qiū) means, literally, knowing from one falling leaf that autumn has come. It names the capacity to infer a larger reality from limited, incomplete signals: to read situational significance that is not explicitly stated. The defining characteristic of this level is judgement under incompleteness: discerning what a fragment means, what matters and what does not, and what the part implies about the whole. Transfer applies a known principle to a new case; contextual inference operates where no clear principle has yet been identified and where the information available is partial, ambiguous, or novel. It depends on situated understanding, a sense built from experience of how a particular domain behaves, and on the judgement to know which signals are significant. Expertise at this level takes the form of recognition-primed judgement, built from experience and exercised under incomplete and time-pressured conditions (Klein, 1998). Much of the judgement that distinguishes a capable graduate from a merely well-

informed one lives here: not in how much they know, but in their ability to interpret incomplete situations and act on them.

3.5 Level 4: Original creation (无中生有)

The classical term 无中生有 (wú zhōng shēng yǒu) means, literally, producing something from nothing. In cognitive terms it names the generation of genuinely novel ideas, frameworks, or works. Such output is not a recombination of existing elements but a true origination. At this level, the learner does not apply, transfer, or infer from what exists; they create what did not previously exist for them. Original creation here does not mean producing something unprecedented in all of human knowledge. Novelty is understood relative to the creator's own starting point: a student who arrives at an argument they were not given, rather than assembling one from templates, is engaged in original creation in the sense that matters educationally. The relevant distinction is between recombination, the skilful rearrangement of existing material, which is itself valuable, and origination, the emergence of something not derivable from its inputs: a new theoretical lens, an original argument, a genuine conceptual leap. That distinction is precisely where, as the next section argues, human cognition and generative AI diverge most sharply.

3.6 The cognitive-depth axis

Taken together, the four levels form an axis of increasing cognitive depth. They are not a sequence of discrete skills to be mastered in turn, nor a ladder a learner climbs and leaves behind; they are cumulative rather than substitutive, each remaining in play throughout expert practice, with higher levels depending on lower ones. Depth here does not rank difficulty, attainment, or value; the axis orders only the degree to which cognitive work can be reduced to rules and patterns rather than requiring situated judgement. Cognitive depth and codifiability move in opposite directions, and it is this property that makes the framework a lens for examining what AI can and cannot do.

4. Mapping AI capability onto the framework

Generative AI models and reproduces patterns in language and data. As cognitive work becomes less reducible to pattern, what happens to a pattern-based system? Examined level by level, capability holds almost completely at Level 1, remains strong but unreliable at Level 2, degrades sharply at Level 3, and lies beyond reach at Level 4.

4.1 Level 1, broad knowledge acquisition and retention: Near-complete capability

At the level of acquiring and reproducing knowledge, generative AI is superhuman. Trained on a corpus no individual could read in a lifetime, it retrieves, summarises, and restates material on demand across domains. One qualification is that the model does not know what it reproduces in the way a person who has retained knowledge does. It predicts likely sequences rather than consulting and verifying a fact, and when likely and true diverge, it produces fluent, confident hallucinations (Huang et al., 2025). Whether this is a temporary engineering shortcoming or a permanent property remains contested in the literature: one line of analysis argues that hallucination is an innate limitation of the architecture rather than a defect that more training will remove (Xu et al., 2024), while the mitigation literature treats it as reducible in practice (Huang et al., 2025). The design implication holds on either reading. Capability at this level is near-complete in coverage but not self-verifying: this is a minor problem when the user can check the output against their own knowledge, but a serious one when they cannot, as is increasingly the case at the higher levels.

4.2 Level 2, analogical transfer: Strong but unreliable

Analogical transfer asks the system to recognise that a principle from one context applies in another. Here generative AI is genuinely strong: it maps structures across domains and produces analogies that are often illuminating, because analogical reasoning is itself largely patterned and the model has seen many worked examples. The limitation is one of basis, not frequency. As Section 3 noted, sound transfer depends on grasping why a principle holds; the model transfers on surface and statistical similarity, not causal

understanding. The result is a capability strong on average but unreliable at the margin. It is unreliable precisely where it matters most: in the unfamiliar or high-stakes case, where a confident but false analogy is hardest to catch.

4.3 Level 3, contextual inference: Where capability degrades

It is at the third level that the relationship between cognitive depth and machine capability changes character. The first two levels involve work that is substantially patterned: a corpus of prior instances supplies regularities to draw on. Contextual inference is different in kind. It is the work of reading a partial, ambiguous, or novel situation and judging what it means: discerning which signals matter, what a fragment implies about a whole that is not visible, and what is about to happen before it has. This is the mode named by the image of the single falling leaf, and the one a pattern-based system is least equipped to perform.

The reason is structural, not a matter of insufficient training. Contextual inference operates on the particular case in its full situational richness, much of which is unstated and absent from any corpus. Consider three cases: the experienced clinician who senses that a patient is more unwell than the chart indicates; the teacher who reads from a classroom's silence that the explanation has not landed; and the analyst who feels a deal is wrong before the numbers say so. Each integrates tacit, contextual cues that were never written down, and could not be, drawing on what Polanyi (1966) called the tacit dimension of knowing, his observation that we know more than we can tell. A system trained on what has been recorded has no access to what was never recorded. It can reproduce the form of expert judgement without performing the judgement itself.

This produces a characteristic and dangerous failure mode. Because the model is fluent at the form of contextual judgement, its output here is often plausible, well-expressed, and wrong in ways that only Level 3 judgement can detect. The user most exposed is the one who has not developed that judgement and so cannot see what the model missed. The Level 1 qualification now bites: at Level 1 the user could check the output against their own knowledge; at Level 3 the only adequate check is the very capability the user may have delegated to the machine.

The risk is therefore pedagogical before it is technical. If a learner routinely offloads contextual inference before developing it, they may never develop it: the concern is not that the machine produces a poor answer, but that the human loses the means of recognising a poor answer. An experimental study found that learners relying on generative AI offloaded regulatory effort in ways that undermined their self-regulated learning (Fan et al., 2025). It is one thing to use a tool that extends a capability one possesses; it is another to use a tool in place of acquiring that capability at all. This distinction between extension and substitution is the hinge on which Section 5 turns. On the structural argument above, the break at Level 3 is not one that more data can close, because the missing input, the unrecorded particulars of a situation, is not the kind of thing training can supply.

4.4 Level 4, original creation: Beyond the boundary

If contextual inference is where capability degrades, original creation is where it is absent. The Section 3 distinction between recombination and origination is decisive. Generative AI is, by construction, a recombinative system: it produces new sequences by sampling from the distribution of what it has seen. Within that space it can be strikingly inventive, generating combinations no human has written. But this is novelty of arrangement, not of origin. It cannot produce what is not derivable in some recombined form from its inputs, because derivation from inputs is the whole of what it does. Generative AI generates, but it does not originate. This neither diminishes recombinative output nor romanticises human creativity, much of which is also recombination. The genuine conceptual leap, the lens that reorganises a field, belongs to a mode that a system defined by its training distribution cannot, on this account, occupy. This claim is a philosophical position rather than an empirical finding: it rests on the conceptual distinction between origination and derivation-from-inputs, not on any measured ceiling of system performance, and it is open to challenge on those terms. On this position, Level 4 is human-exclusive not because machines are not yet good enough, but because origination and derivation are different acts.

4.5 The capability gradient

Across the four levels, the capability of a pattern-based system is the inverse of cognitive depth: highest where work is most codifiable, and falling away as work comes to require situated judgement and genuine origination. Any undifferentiated stance treats as uniform a capability that is in fact steeply graded. That is the argument of Section 5.

5. Implications for AI integration in higher education

This section translates the capability gradient into a design heuristic and illustrates its use through two anonymised cases from undergraduate business education.

5.1 Differentiated integration: The principle

The principle follows from the inverse relationship between cognitive depth and machine capability. At Level 1, AI can soundly support much of the work: retrieving, collating, and summarising knowledge, with the learner verifying rather than passively accepting. At Level 2, AI can assist, but its analogies require checking against causal understanding the learner must supply. At Level 3, the relationship reverses: the work of judging incomplete situations is precisely what learners must do themselves, both because AI performs it unreliably and, more fundamentally, because offloading it prevents the capability from forming. Protecting this work aligns with calls to foreground students' evaluative judgement, their capacity to judge the quality of work, as a distinctly human capability at a time of generative AI (Bearman et al., 2024). At Level 4, the work is the learner's alone; AI may serve as material or stimulus, but origination cannot be delegated to a system that recombines.

Stated as a design heuristic: the deeper the cognitive work a task is meant to develop, the more the design must protect that work from substitution, and the more legitimate it becomes to delegate the shallower work that surrounds it, within a single task as well as across tasks. The framework serves as a diagnostic lens through which an educator reads a task, locates the cognitive work that constitutes its learning purpose, and assigns AI a role consistent with that purpose. Its most immediate use is therefore with educators rather than students: it gives a course team a shared vocabulary for saying which parts of a task AI may support and why, replacing a blanket stance on the whole task with a reading of its parts.

5.2 A worked example: One assessment, four levels

In a team-based experiential assessment in an undergraduate management course, students participate in a multi-round group simulation in which teams make sequential decisions under pressure, and each student then writes an individual analysis of their team's dynamics, applying course frameworks to what they experienced. The task is a single assessment with one rubric and one submission. Viewed through the framework, the assessment spans all four levels of cognitive work, and AI's appropriate role differs at each.

At Level 1, the task requires command of course concepts and supporting literature: teamwork theories, decision-making models, and the academic sources that ground the analysis. Retrieving relevant theory, summarising sources, checking definitions: this is work AI performs at near-complete capability, and delegating parts of it may be appropriate, provided the learner verifies what is retrieved against authoritative sources, since fluent fabrication remains a Level 1 failure mode.

At Level 2, the task requires applying general frameworks to the specific simulation: recognising, for example, that a model of group decision-making maps onto what the team did in round three. AI can propose such mappings, and they are often useful starting points, but because its transfer rests on surface similarity rather than causal understanding, the learner must own the verification, asking whether the mechanism the framework describes was actually the mechanism at work.

At Level 3, the character of the task changes. The heart of the assessment is the student's reading of their own team: judging why a particular exchange shifted the group's direction, what an unspoken tension meant, and which of the many events of the simulation were significant and which were noise. This is contextual inference

performed on an experience only the student had. The situational particulars on which the judgement depends, the tone of a disagreement, the hesitation before a decision, the difference between what was said and what was meant, were never recorded and exist in no corpus. AI brought to this work can produce only a fluent simulation of team analysis: generic dynamics described in plausible language. A student who delegates the work receives, in place of their own developing judgement, an articulate account of a team that never existed. The design implication is protective: this is the cognitive work the assessment exists to develop, and the design must require that it be done by the student, from their own experience and in their own voice.

The same logic extends beyond the written submission to the discussion that surrounds it. If students arrive at team discussion with generated positions rather than worked-through interpretations, the live exchange thins: a student can state a view but cannot extend, defend, or revise it when the conversation moves, because the position was never theirs to begin with. This is precisely the pattern the framework predicts: offloading Level 3 work does not merely risk a poorer artefact; it weakens the conditions under which live, situated judgement is formed. There is also a disciplinary point that the framework brings out. Management has always meant accountability for cognitive work one did not perform oneself, and generative AI simply makes that condition universal, which is why protecting the student's own Level 3 judgement is not peripheral to a management course but central to it. At Level 4, finally, the assessment invites an original synthesis: an insight about the team that no course framework supplied and that could not be assembled from templates. AI may serve as stimulus here, but the synthesis is reserved to the learner. Table 2 summarises the reading.

Table 2.
One assessment read through the framework: the appropriate role of AI at each level

Level	Cognitive work in this assessment	Appropriate AI role	Design treatment
1	Theories, models, and sources	Retrieve, summarise, check definitions	Permit with source verification
2	Mapping frameworks to simulation events	Propose candidate mappings	Permit as starting points; learner owns causal verification
3	Interpreting team dynamics from lived experience	No generative substitution; reflective prompts only	Protect: require the student's own interpretation and voice
4	Original synthesis and insight	Stimulus at most	Reserve wholly to the learner

A task-based rule would permit or prohibit AI for the whole assessment, and either choice is wrong for half of it. The framework yields a differentiated reading instead: permit and verify at Levels 1 and 2, protect at Level 3, reserve at Level 4. One assessment, one coherent policy, from a small set of stable questions rather than a per-task ruling.

5.3 Embedding the differentiation: A scaffolding case

The second case, from a foundational disciplinary course in undergraduate business education, illustrates one mechanism. The course embeds a scaffolding module built on the Co-intelligence Activity Chain (CAC) (Nikolic, 2026), a model extending Mollick's (2024) notion of human-AI co-intelligence into a teachable sequence that structures AI-assisted work. The student establishes their own thinking before engaging AI, then engages it actively and critically, then returns to their own judgement to evaluate and integrate what the engagement produced. In the course's major written assessment, this is operationalised as a scored rubric criterion: students earn marks for demonstrating that they applied the sequence, or for explaining a reasoned decision not to use AI at all.

Reading through the framework, the design's logic is visible. The CAC sequence brackets AI engagement, which operates mainly at Levels 1 and 2, between two phases of the student's own cognition where Level 3 judgement is exercised: deciding what to ask, evaluating what came back, and judging what it is worth. The scored criterion converts the differentiation from advice into assessment: rather than an unenforceable list of permitted and prohibited behaviours a single criterion rewards the cognitive sequence that keeps judgement with the learner. The framework did not generate this design, which predates it; it explains why the design works and offers the means of adapting the same logic to courses whose Level 3 work differs entirely.

6. Discussion and limitations

The contribution of this paper is conceptual, and its claims should be read at that level. As noted in Section 1, the work involved no data collection and no human participants, and ethics approval was not required. The framework's classical Chinese origins are offered as a source, not as a claim in intellectual history; the paper makes no philological argument about the four concepts, and the framework must stand or fall on its contemporary analytical value. Three further limitations define the boundary of what is being asserted.

The capability mapping in Section 4 concerns current and near-term generative AI. The structural argument, that contextual inference depends on situational particulars that no corpus records, gives reason to expect the Level 3 boundary to be durable, but the pace of development means the mapping should be revisited as systems change rather than treated as a fixed prediction. The Level 4 claim, as Section 4.4 makes explicit, is a philosophical position rather than an empirical finding.

The framework is not a measurement taxonomy. It does not grade student work, sequence curriculum, or replace outcome frameworks such as Bloom's taxonomy or SOLO in the work they were designed to do. It is a diagnostic lens for design judgement: a way of locating the cognitive work that a task exists to develop and of assigning AI a role consistent with that purpose. The two kinds of framework are complements, not rivals.

The cases in Section 5 are illustrative, not evaluative. They describe real, anonymised course designs; they do not demonstrate improved learning outcomes. Empirical work is the natural next step, and two questions are directly testable: whether designs that protect Level 3 work produce measurably stronger contextual judgement than designs that do not, and whether educators applying the framework converge in their readings of the same task. The framework is now informing AI-integration design across the core courses of an undergraduate business program, and its empirical evaluation in that setting is planned as a subsequent study.

7. Conclusion

This paper began from a mismatch. Higher education's rules for generative AI are organised by task, while the thing they are meant to protect, learning, is organised by cognitive work. The four-level framework addresses that mismatch directly: by ordering cognitive work from the codifiable to the situated, it makes visible what task-level guidance cannot, that AI capability is the inverse of cognitive depth. The practical consequence is differentiation. Within a single assessment, AI can soundly support the work of acquisition and transfer while the work of contextual inference and original creation is deliberately protected, and a small set of stable cognitive questions can replace an unbounded list of task rules.

If machines increasingly supply what can be codified, education must prioritise what cannot: judgement under incomplete conditions and learner-owned origination. Rebalancing learners and algorithms is therefore not a matter of rationing a tool but of knowing which cognition to keep human.

It is fitting that this lens emerged from outside the analytic tradition that shaped educational technology. By grasping modes of knowing as wholes, the four classical concepts help bring the edge of the machine into view.

References

Anderson, L. W., & Krathwohl, D. R. (Eds.). (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives* (Complete ed.). Longman.

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

Biggs, J. (1996). Enhancing teaching through constructive alignment. *Higher Education, 32*(3), 347–364. https://doi.org/10.1007/BF00138871

Biggs, J. B., & Collis, K. F. (1982). *Evaluating the quality of learning: The SOLO taxonomy (Structure of the Observed Learning Outcome)*. Academic Press.

Bloom, B. S., Engelhart, M. D., Furst, E. J., Hill, W. H., & Krathwohl, D. R. (1956). *Taxonomy of educational objectives: The classification of educational goals. Handbook I: Cognitive domain.* David McKay.

Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education, 20*, Article 38. https://doi.org/10.1186/s41239-023-00408-3

Corbin, T., Dawson, P., Nicola-Richmond, K., & Partridge, H. (2025). "Where's the line? It's an absurd line:" Towards a framework for acceptable uses of AI in assessment. *Assessment & Evaluation in Higher Education, 50*(5), 705–717. https://doi.org/10.1080/02602938.2025.2456207

Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education, 20*, Article 22. https://doi.org/10.1186/s41239-023-00392-8

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

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7*(2), 155–170. https://doi.org/10.1207/s15516709cog0702_3

Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies, 15*(1), Article 6. https://doi.org/10.3390/soc15010006

Gonsalves, C. (2026). Generative AI's impact on critical thinking: Revisiting Bloom's taxonomy. *Journal of Marketing Education, 48*(1), 4–19. https://doi.org/10.1177/02734753241305980

Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., & Liu, T. (2025). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems, 43*(2), Article 42. https://doi.org/10.1145/3703155

Klein, G. A. (1998). *Sources of power: How people make decisions*. MIT Press.

Li, J. (2012). *Cultural foundations of learning: East and West*. Cambridge University Press. https://doi.org/10.1017/CBO9781139028400

Lodge, J. M., Howard, S., Bearman, M., & Dawson, P. (2023). *Assessment reform for the age of artificial intelligence*. Tertiary Education Quality and Standards Agency. https://www.teqsa.gov.au/guides-resources/resources/corporate-publications/assessment-reform-age-artificial-intelligence

Mollick, E. (2024). *Co-intelligence: Living and working with AI*. Portfolio.

Nikolic, S. (2026, April 1). *CAC co-intelligence: Teaching AI use that strengthens the mind & supports GenAI integration*. https://www.sashanikolic.com/post/cac-co-intelligence-teaching-ai-use-that-strengthens-the-mind

Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2001). Culture and systems of thought: Holistic versus analytic cognition. *Psychological Review, 108*(2), 291–310. https://doi.org/10.1037/0033-295X.108.2.291

Polanyi, M. (1966). *The tacit dimension*. Doubleday.

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

Xia, Q., Weng, X., Ouyang, F., Lin, T. J., & Chiu, T. K. F. (2024). A scoping review on how generative artificial intelligence transforms assessment in higher education. *International Journal of Educational Technology in Higher Education, 21*, Article 40. https://doi.org/10.1186/s41239-024-00468-z

Xiao, P., Chen, Y., & Bao, W. (2023). *Waiting, banning, and embracing: An empirical analysis of adapting policies for generative AI in higher education* [Preprint]. *arXiv*. https://doi.org/10.48550/arXiv.2305.18617

Xu, Z., Jain, S., & Kankanhalli, M. (2024). *Hallucination is inevitable: An innate limitation of large language models* [Preprint]. *arXiv*. https://doi.org/10.48550/arXiv.2401.11817

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

ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

Lin, M. (2026). What AI can and cannot do: A cognitive-depth framework for differentiated AI integration in higher education. In Cochrane, T., Narayan, V., Poole, T., Lai, P., Lake, P., Moore, O., & Li, L. (Eds.), *The Future is Connection: Collaboration, Creativity and Compassion in a Digital Society*. Proceedings ASCILITE 2026. Brisbane (pp. 293–303). https://doi.org/10.65106/apubs.2026.3267

Note: All published papers are refereed, having undergone a double-blind peer-review process.
The author(s) assign a Creative Commons by attribution license enabling others to distribute, remix, tweak, and build upon their work, even commercially, as long as credit is given to the author(s) for the original creation.

© Lin, M. 2026