Bloom's Taxonomy in the Age of Generative AI: From Creation To Responsible Judgement Through Verify, Orchestrate, and Reflect (VOR)

Prof. Rabee Rustum

Institute of Sustainable Built Environment, Heriot-Watt University, Dubai Campus, UAE.

DOI: https://doi.org/10.47772/IJRISS.2026.101100082

Received: 28 October 2026; Accepted: 03 November 2026; Published: 21 November 2026

ABSTRACT

The rapid adoption of generative artificial intelligence (AI) in higher education is challenging established assumptions about the relationship between student-produced artefacts and evidence of cognitive achievement. This perspective paper argues that these developments do not render Bloom's Taxonomy dated, but require its reinterpretation for learning environments in which cognitive work can be distributed between learners and AI systems. The perspective retains Bloom's six cognitive processes of Remember, Understand, Apply, Analyse, Evaluate, and Create while reconsidering how each is demonstrated when AI can retrieve information, generate content, analyse problems, and contribute to creative production. AI literacy is positioned as an enabling capability that operates across all levels rather than as an additional cognitive category. The paper further proposes Verify, Orchestrate, and Reflect (VOR) as three cross-cutting capabilities for AI-mediated learning. Verify concerns epistemic judgement and the ability to establish the credibility and evidential basis of information; Orchestrate concerns the strategic allocation, coordination, and supervision of cognitive work across human and AI resources; and Reflect concerns the ability to examine, explain, and take ownership of the reasoning and decisions underpinning an outcome. Their integration supports responsible judgement, conceptualised not as a seventh level of Bloom's hierarchy but as an emergent quality of higher-order performance. The paper therefore proposes a shift from production-centred evidence of learning to judgement-centred evidence, arguing that creation remains important but its educational value increasingly depends on the quality of human judgement surrounding it. The paper develops implications for assessment through oral defence, authentic tasks, process documentation, AI critique, and verification-focused activities. The proposed framework provides a basis for curriculum and assessment design that preserves Bloom's enduring emphasis on cognitive complexity while responding to the epistemic, strategic, and reflective demands of human-AI learning environments.

Keywords: Bloom's Taxonomy; generative artificial intelligence; higher education; human-AI collaboration; AI literacy; assessment; verification; orchestration; reflection; judgement.

INTRODUCTION

Generative artificial intelligence (GenAI) has altered a basic assumption on which much higher-education assessment has historically rested: a sophisticated academic artefact is principally the product of the learner who submits it. Contemporary systems can generate essays, code, summaries, designs, and problem solutions that resemble competent student work, making the relationship between observable output and underlying cognition less direct (O'Dea, 2024; Xia et al., 2024; Stephenson & Armstrong, 2026). The resulting challenge concerns not only academic integrity, but assessment validity: what evidence allows educators to infer what a learner understands, can do, and can responsibly justify when AI can participate in production?

This problem has renewed interest in Bloom's Taxonomy. Bloom's framework remains deeply embedded in curriculum design, learning outcomes, and assessment, yet generative AI challenges conventional ways of operationalising its higher levels. In particular, when AI can contribute substantially to analysis and creation, the sophistication of a submitted product cannot, by itself, establish that the learner performed equivalent cognitive

work (Kasneci et al., 2023; Xia et al., 2024; Lubbe et al., 2025). At the same time, AI-rich learning increases the importance of critical evaluation, information verification, AI literacy, human oversight, and accountable decision-making (Miao and Holmes, 2023; Hazari, 2024; Bacalso et al., 2026).

The gap addressed here is therefore not whether Bloom's six cognitive processes remain useful, but whether conventional interpretations are sufficient when cognition and production are increasingly distributed across human and technological resources. Existing responses to generative AI often emphasise reassessing Bloom, AI literacy, critical thinking, or assessment redesign. Less explicit is an integrated account of three capabilities needed when AI participates in cognitive work: establishing whether information can be trusted, strategically allocating and coordinating work between human and AI resources, and making the learner's reasoning and ownership visible.

This perspective paper makes three claims. First, Bloom's cognitive processes remain useful and should be reinterpreted rather than replaced. Second, AI literacy should be treated as an enabling capability across those processes rather than as an additional rung in the taxonomy. Third, the paper proposes Verify, Orchestrate and Reflect (VOR) as cross-cutting capabilities whose integration supports responsible judgement. The resulting model shifts the evidential question from 'What can the student produce?' toward 'What can the student understand, verify, coordinate, justify and take responsibility for when AI can participate in producing an answer?' The paper then considers implications for assessment and outlines a research agenda for testing the framework.

From Creation to Judgement" does not imply replacing Create as the apex of Bloom's hierarchy. Rather, it describes a shift in the evidential basis of higher-order achievement, from the sophistication of what is produced towards the quality of judgement surrounding its production and use.

WHY BLOOM STILL MATTERS IN THE AI ERA

2.1 Original Purpose and Enduring Value

Bloom's Taxonomy was developed as a language for describing educational objectives and differing forms of cognitive engagement, not as a catalogue of technologies or assignment types (Bloom et al., 1956). Krathwohl (2002) recast the cognitive categories as active processes and reinforced their use in curriculum and assessment design. Its enduring value therefore lies in helping educators ask what kind of thinking a learning outcome requires, rather than prescribing the medium through which that thinking must be demonstrated.

This distinction is especially important in AI-mediated education. Technologies can alter access to information and the production of academic artefacts without eliminating the need to understand concepts, apply knowledge, analyse relationships, evaluate evidence or construct solutions (UNESCO, 2023; O'Dea, 2024). Bloom is therefore most defensible as a heuristic for cognitive demand rather than as a strict account of how learning unfolds. Cognitive processes can overlap, recur, and combine differently across tasks and disciplines; creation can require evaluation, while evaluation can depend on remembering and understanding.

2.2 Enduring Strengths

Three strengths justify retaining Bloom. First, it provides a durable vocabulary for cognitive complexity and higher-order thinking. Second, it supports alignment by helping educators articulate the cognitive demands of intended learning outcomes and assessments. However, constructive alignment is more properly associated with the wider curriculum-design tradition than with Bloom alone. Third, the taxonomy has proved adaptable: the Krathwohl (2002) revision shows that the framework can evolve while retaining its central focus on intellectual activity (Krathwohl, 2002). Recent work on AI and assessment similarly supports adapting educational frameworks rather than replacing them wholesale (Xia et al., 2024; Ilieva et al., 2025; Lubbe et al., 2025).

Generative AI arguably increases the importance of distinguishing levels and qualities of cognitive engagement. When fluent outputs are easy to generate, educators need stronger ways to distinguish retrieval from understanding, plausible argument from defensible judgement, and artefact production from intellectual

ownership. Bloom continues to provide a useful starting structure for these distinctions, provided the framework is not treated as a simple linear hierarchy or used to infer cognition automatically from a product's sophistication (Miao and Holmes, 2023; Bacalso et al., 2026).

2.3 Limitations of Conventional Applications in AI-Rich Environments

The central limitation generative AI exposes lies partly in conventional applications of Bloom rather than in the taxonomy's cognitive categories themselves. Educational practice has often inferred higher-order cognition from observable products such as essays, designs, projects, and code. Generative AI weakens that inference because sophisticated products may now reflect very different combinations of learner effort, prior knowledge, and machine contribution (O'Dea, 2024; Xia et al., 2024; Stephenson & Armstrong, 2026). Product quality therefore remains relevant, but it is no longer sufficient evidence of the cognitive process that produced it.

A second limitation is that Bloom offers little explicit guidance for allocating cognitive work across external resources. Learners now decide not only how to solve a problem but whether AI should be used, which system or resource is appropriate, what context should be supplied, when outputs should be rejected, and how machine-generated contributions should be combined with disciplinary expertise (Jin et al., 2024; Ayyoub et al., 2025). These decisions are central to human-AI problem solving but are not clearly represented in conventional Bloom mappings.

A third limitation concerns epistemic verification. Generative AI can produce fluent but inaccurate claims, fabricated references, and biased or unsupported conclusions. Analysis and evaluation can encompass aspects of these tasks, but AI makes establishing whether a claim is true, adequately evidenced, and trustworthy sufficiently important to merit explicit conceptual visibility (Bobula, 2024; UNESCO, 2023; Bacalso et al., 2026).

Finally, AI complicates Create's position as a readily observable high of achievement. Creation remains cognitively demanding, but AI can participate directly in producing the artefact. Decisive evidence increasingly lies in how the learner frames the problem, evaluates alternatives, directs or constrains AI, modifies outputs, and takes responsibility for the result (Xia et al., 2024; Lubbe et al., 2025). Similarly, reflection is not a new educational construct, but AI increases its evidential importance because learners may need to demonstrate how and why a particular output emerged from a human-AI process.

These limitations do not require replacing Bloom. They require two complementary moves: reinterpret the six cognitive processes for AI-mediated learning, and make explicit the cross-cutting capabilities required when cognitive work is shared with intelligent systems.

REINTERPRETING BLOOM FOR AI-MEDIATED LEARNING

The proposed reinterpretation retains Bloom's six cognitive processes while changing how competence at each level is evidenced. It does not assume a rigid sequence. Rather, the six processes provide a language for cognitive demand, while AI literacy enables learners to engage appropriately with AI across levels (Miao and Holmes, 2023; Hazari, 2024; Ayyoub et al., 2025).

Table 1. Reinterpreting Bloom's cognitive processes for AI-mediated learning

Bloom process	AI-era interpretation	Illustrative evidence of learning
Remember	Retrieve and recognise relevant disciplinary knowledge, including knowing when external retrieval is needed.	Accurate identification of concepts, terms, sources and relevant prior knowledge.
Understand	Explain, contextualise, transfer and defend conceptual meaning independently of AI-generated fluency.	Explanation in the learner's own reasoning; transfer to a new context; correction of misconceptions.
Apply	Use disciplinary knowledge independently and with appropriate AI support.	Selection and adaptation of methods or tools to address an authentic task.

Analyse	Decompose problems and interrogate human- and AI-generated reasoning, assumptions and relationships.	Identification of assumptions, inconsistencies, biases, dependencies and missing information.
Evaluate	Judge evidence, alternatives and recommendations against disciplinary, contextual and ethical criteria.	Defensible selection among alternatives using explicit criteria and evidence.
Create	Generate, integrate and refine solutions, including through responsible human-AI collaboration.	Original synthesis plus transparent direction, refinement and ownership of the outcome.

This reinterpretation deliberately avoids placing 'prompting' or AI itself within the hierarchy. AI literacy is better understood as an enabling capability that affects how learners perform cognitive work at every level. Moreover, verification is not confined to Remember, orchestration is not confined to Apply or Create, and reflection is not an additional final rung. These capabilities can appear wherever a task requires them. This distinction sets the foundation for the Verify, Orchestrate and Reflect (VOR) Framework developed below.

4. Beyond Bloom: Verify, Orchestrate and Reflect

Bloom describes forms of cognitive engagement, but AI-mediated learning also requires capabilities concerned with the learner's relationship to the cognitive resources being used. This paper proposes three such cross-cutting capabilities: Verify, Orchestrate and Reflect. They are not substitutes for Bloom's six processes and are not claimed to be new. Their contribution is to make explicit three functions whose educational significance increases when AI can retrieve, generate, and recommend on behalf of a learner.

4.1 Verify: Epistemic Judgement

Verify is defined here as the capability to establish whether a claim, source, or output is sufficiently accurate, supported, and trustworthy for the purpose at hand. It differs from Analyse and Evaluate in emphasis: Analyse asks how a claim or solution is structured; Evaluate asks how good, appropriate, or valuable it is; Verify asks whether the relevant factual and evidential claims can be warranted. Verification therefore operationalises epistemic judgement in circumstances where fluency is not a reliable indicator of truth (Miao and Holmes, 2023; Bobula, 2024).

- Is the claim factually correct and current?
- What evidence supports it, and can that evidence be authenticated?
- Are citations genuine and appropriately represented?
- What assumptions, omissions, or biases affect confidence in the output?
- What independent source or method can corroborate the claim?

Verification has disciplinary forms: an engineer may recalculate an AI-derived result, a researcher may authenticate a reference, and a business student may triangulate a market claim. The common educational capacity is the ability to justify why information should be trusted before using it in reasoning or action. In this sense, verification is both an AI-literacy capability and an expression of higher-order judgement (Hazari, 2024; Bacalso et al., 2026).

4.2 Orchestrate: Strategic Human-AI Coordination

Orchestrate is defined as the strategic allocation, sequencing, coordination, and supervision of cognitive work across human and AI resources to pursue a defined goal. It therefore goes beyond prompt engineering. Prompting concerns interaction with a system; orchestration concerns the architecture of the whole problem-solving process (Jin et al., 2024; Miao and Holmes, 2023).

- Allocation: decide which parts of a task should be performed by the learner, other people, conventional tools, or AI.
- Selection: choose an appropriate resource or AI system for the task and its risks.

- Coordination: sequence and integrate contributions from different resources.
- Supervision: monitor outputs and decide when to revise, override, escalate, or stop AI use.

Orchestration is therefore a form of AI-enabled problem solving in which the learner remains responsible for configuring and supervising a distributed work process. Its educational value lies not in using more AI, but in making proportionate and defensible decisions about when, where, and how AI should contribute (Ayyoub et al., 2025; Bacalso et al., 2026).

4.3 Reflect: Making Thinking and Ownership Visible

Reflect is defined as the capability to examine and communicate the reasoning, assumptions, decisions, and revisions through which an outcome was reached. Reflection is well established in education and metacognition; the claim here is not that AI creates reflection, but that AI increases its importance as evidence of learning and intellectual ownership. When output generation can be distributed between learner and machine, explanation of the process can reveal cognitive engagement that the final artefact cannot show (Hazari, 2024; Xia et al., 2024).

- Why was a particular AI contribution accepted, rejected, or modified?
- What evidence or disciplinary principle informed the decision?
- How did the learner's understanding change during the task?
- What limitations remain in the final solution?
- For which parts of the outcome does the learner accept responsibility?

Reflection therefore supports metacognition and self-regulation while also making judgement visible for assessment. It can be evidenced through oral defence, reflective commentary, decision logs, portfolios and annotated revisions (Miao and Holmes, 2023; Lubbe et al., 2025).

4.4 The Integrated VOR Framework

VOR is proposed as a cross-cutting layer rather than an extension of Bloom's hierarchy. Verify can occur during remembering, analysis, evaluation, or creation; Orchestrate is especially visible during application and creation but can shape the whole task; Reflect can occur before, during, and after any cognitive activity. AI literacy provides the enabling knowledge and skills needed to work with AI. At the same time, responsible Judgement is the integrative outcome that emerges when learners verify claims, orchestrate resources, and reflect on decisions.

Table 2. Relationship among Bloom, VOR, AI literacy and responsible judgement

Element	Role in the model	Primary question
Bloom's six processes	Describe the cognitive demand of learning activities.	What kind of thinking is required?
AI literacy	Knowledge and skills that enable informed interaction with AI.	What do I need to understand about AI to use it appropriately?
Verify	Establish credibility, accuracy, and evidential warrant.	What should I trust, and why?
Orchestrate	Allocate, coordinate, and supervise human and AI resources.	Who or what should do which part, and how should contributions be integrated?
Reflect	Examine and communicate reasoning, revision, and ownership.	Why did I make these decisions, and what did I learn?
Responsible judgement	Integrate evidence, context, values, accountability, and VOR capabilities.	What should be done, and can I defend and take responsibility for that decision?

Figure 1 illustrates the relationship between Bloom’s Taxonomy, AI literacy, and the proposed Verify–Orchestrate–Reflect (VOR) capabilities. Bloom’s six cognitive processes remain unchanged and form the framework's foundation. VOR operates across all levels rather than constituting additional stages within the

taxonomy, while AI literacy enables effective engagement with AI throughout the learning process. The integration of these elements supports the development of Responsible Judgement, which is conceptualised as a central outcome of learning in the age of generative AI.

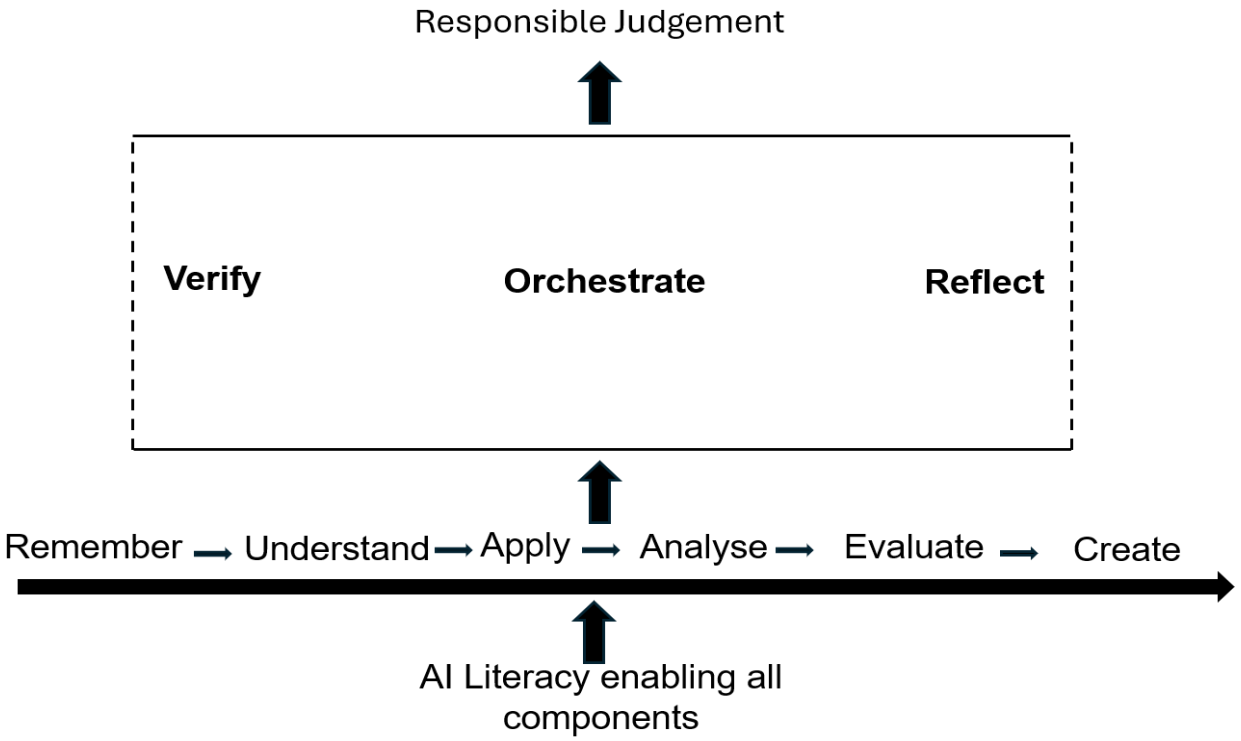


Figure 1. Integrated Bloom–VOR Framework

5. From Creation to Judgement

The framework's theoretical contribution can be stated in three propositions. First, generative AI weakens the correspondence between artefact sophistication and learner cognition. A polished essay, design, or program may still show achievement, but its quality alone no longer indicates how much relevant cognitive work the learner performed (O'Dea, 2024; Xia et al., 2024). Assessment must therefore seek additional evidence of the cognitive and metacognitive processes surrounding production.

Second, AI makes cognitive production more distributed. Learners can allocate parts of retrieval, drafting, comparison, synthesis, and revision to technological systems. The educational responsibility retained by the learner therefore becomes especially visible in decisions about problem framing, evidence, tool use, accepting or rejecting outputs, and integrating contributions into a defensible solution (Jin et al., 2024; Miao and Holmes, 2023). The idea of distributed cognitive work is descriptive: not to claim that cognition is reducible to the human-AI system, but to recognise that observable academic production may emerge from coordinated contributions rather than from an individual acting alone.

Third, these changes imply a shift from production-centred evidence toward judgement-centred evidence. The claim is not that judgement replaces creativity or that Create should be removed from Bloom. Rather, when production is partly automatable, the educational significance of creation increasingly depends on the quality of judgement surrounding it. A learner who can verify evidence, compare alternatives, coordinate resources, explain revisions, and accept responsibility demonstrates forms of intellectual agency that cannot be inferred from the artefact alone (Miao and Holmes, 2023; Lubbe et al., 2025).

This distinction also clarifies intellectual ownership. In AI-mediated work, ownership is not exhausted by sole authorship. It can be demonstrated through purposeful direction, critical modification, transparent attribution, and accountable acceptance of the final outcome. Responsible judgement therefore integrates epistemic,

strategic, reflective and ethical dimensions. It is not proposed as a seventh Bloom level but as an integrative quality of higher-order performance that may be expressed especially strongly in evaluation and creation.

The philosophical implication is a change in the question higher education asks of student work. The conventional question "Can the student produce the answer?" becomes insufficient. A more defensible question is: "Can the student frame the problem, determine what to trust, coordinate resources, explain the reasoning, and take responsibility for the outcome?" This reframing preserves Bloom's concern with cognitive complexity while recognising that the relationship between cognition, agency and production has changed in human-AI learning environments.

Judgement, as conceptualised here, is broader than Bloom's Evaluate. Evaluate describes a cognitive process of making determinations against criteria and standards. Responsible judgement refers to integrating epistemic, strategic, reflective, contextual, and ethical considerations when deciding what to believe, what action to take, how AI should contribute, and for which outcomes the learner accepts responsibility. Evaluation is therefore one cognitive process contributing to judgement, rather than being synonymous with it.

IMPLICATIONS FOR ASSESSMENT

If artefacts provide less complete evidence of learner cognition, assessment should triangulate product evidence with evidence of reasoning, process and accountability. The goal is not to eliminate AI use or abandon conventional assessment, but to design combinations of methods that make relevant cognitive work visible and align with intended learning outcomes (Xia et al., 2024; Ilieva et al., 2025).

6.1 Viva and Oral Defence

A viva or oral defence can provide additional evidence of understanding and ownership by requiring learners to explain concepts, defend methods, justify use of AI, respond to challenge and acknowledge limitations. It should not be treated as definitive proof of authenticity, because explanations can be rehearsed, but as one source of evidence within an assessment design. Used proportionately, oral defence is especially suited to Reflect and Verify because it exposes the learner's reasoning and ability to defend claims (Xia et al., 2024; Lubbe et al., 2025).

6.2 Authentic and Contextualised Tasks

Authentic tasks can require learners to work with contextual constraints, incomplete evidence, stakeholder needs and disciplinary standards rather than answer generic prompts. Case studies, consultancy tasks, simulations and workplace problems can therefore assess application and orchestration in ways that make contextual judgement important. The point is not that AI cannot support such tasks, but that the appropriateness of AI-generated proposals cannot be inferred from output alone and must be judged against the specific context (Miao and Holmes, 2023; Ilieva et al., 2025).

6.3 Process Documentation and Portfolios

Process documentation can reveal how ideas, evidence and decisions developed over time. Drafts, revision histories, decision logs, annotated AI interactions and reflective portfolios can make verification, orchestration and reflection more visible. These records should be sampled and purposefully designed rather than used for exhaustive surveillance: logs can be fabricated, documentation can add workload, and reliable criteria are needed to assess process evidence. Used carefully, they complement rather than replace evaluation of the final product (Hazari, 2024; Bacalso et al., 2026).

6.4 AI Critique and Comparative Evaluation

AI critique tasks ask learners to judge the quality of an AI response. They may compare human and AI solutions, identify weaknesses in reasoning, examine bias or omissions, and recommend improvements. Their primary

purpose is evaluative: to reveal whether a learner can apply disciplinary criteria and critical judgement to machine-generated content (Bobula, 2024; Lubbe et al., 2025).

6.5 Verification-Focused Assessment

Verification tasks relate to, but differ from, AI critique. Whereas critique asks how good or appropriate an AI response is, verification asks whether specific factual and evidential claims are adequately supported. Activities may include authenticating references, triangulating claims, reproducing calculations, checking provenance, or defending the evidential basis for a decision. Verification-focused assessment therefore operationalises epistemic judgement directly and can be adapted to disciplinary standards (Miao and Holmes, 2023; Ayyoub et al., 2025).

Table 3 summarises the assessment methods aligned with the VOR capabilities. The table shows that no single method is sufficient. A stronger assessment strategy combines approaches in ways appropriate to the intended learning outcomes, disciplinary context, stakes and available resources. The central design principle is triangulation: when AI can contribute substantially to products, educators should seek complementary evidence of understanding, verification, strategic coordination, reflection and accountable judgement.

Table 3. Assessment methods aligned with the VOR capabilities

Assessment method	Primary VOR emphasis	Illustrative evidence of learning
Viva / oral defence	Reflect; Verify	Explains reasoning, defends claims, justifies AI use and acknowledges limitations.
Authentic / contextualised task	Orchestrate; Reflect	Coordinates resources and applies judgement under real contextual constraints.
Process documentation/ portfolio	Reflect; Orchestrate	Shows decision pathways, revisions, AI contribution, and evolution of thinking.
AI critique / comparative evaluation	Verify; Reflect	Evaluates quality, bias, reasoning, and appropriateness of alternative outputs.
Verification-focused task	Verify	Authenticates sources and claims, triangulates evidence, and demonstrates evidential warrant.

LIMITATIONS AND FUTURE RESEARCH

The proposed Verify-Orchestrate-Reflect (VOR) framework should be regarded as a conceptual model intended to stimulate discussion and guide future empirical investigation rather than as a validated educational framework. While the paper advances a theoretical reinterpretation of Bloom's Taxonomy for AI-mediated learning, the framework has not yet been empirically tested, and its effectiveness in improving learning, assessment validity, or student outcomes remains to be established.

A second limitation concerns conceptual boundaries. Although Verify, Orchestrate, and Reflect are proposed as distinct capabilities, they share characteristics with established constructs such as critical thinking, metacognition, self-regulated learning, information literacy, digital literacy, and AI literacy. Further theoretical development is therefore needed to clarify whether VOR represents a novel construct, a reconfiguration of existing constructs, or an integrative framework that brings these capabilities together within the context of AI-mediated learning. Clearer conceptual differentiation is particularly important for Verify, which overlaps with critical evaluation and information literacy, and Reflect, which shares features with established metacognitive and reflective-learning traditions.

Third, the framework currently lacks validated measures and assessment criteria. Although this paper outlines illustrative assessment approaches, it does not yet specify measurable indicators, developmental levels, or scoring criteria for the three VOR capabilities. Consequently, it remains unclear how educators might reliably assess students' ability to verify AI-generated information, orchestrate human-AI collaboration, or reflect upon their decision-making processes across disciplines and educational contexts. The practical value of the

framework will therefore depend on the development of robust assessment rubrics and observable performance indicators.

A fourth limitation relates to context and scope. The discussion is primarily situated within higher education, reflecting both the origins of Bloom's Taxonomy and the current concentration of generative AI adoption in universities. Consequently, the applicability of the framework to school education, vocational training, professional development, and lifelong learning contexts remains uncertain. Differences in learner maturity, curriculum requirements, levels of autonomy, and access to AI technologies may require significant adaptation of the framework across educational sectors.

Finally, implementation challenges require careful consideration. Access to AI tools, digital infrastructure, institutional support, and prior technological experience vary substantially among students and educational settings. VOR-informed assessment approaches that rely on AI interaction, oral defence, process documentation, or reflective evidence may unintentionally disadvantage learners with limited technological access, lower levels of AI competence, language barriers, or disabilities. Furthermore, AI systems themselves may introduce bias, inaccuracies, cultural limitations, and accessibility concerns that affect the fairness, validity, and inclusiveness of assessment practices. These factors should be considered before widespread implementation of the framework.

Future research should therefore proceed along several complementary directions. First, empirical studies are needed to test and validate the VOR framework across different disciplines, educational levels, and learning environments. Quantitative, qualitative, and mixed-methods studies could examine whether Verify, Orchestrate, and Reflect can be reliably identified, measured, and distinguished from related constructs such as critical thinking, metacognition, and AI literacy.

Second, researchers should develop and validate assessment rubrics, performance descriptors, and measurable indicators for each VOR capability. Establishing construct validity, reliability, and developmental progression levels will be essential if the framework is to support curriculum design and assessment practice.

Third, classroom-based implementation studies should investigate how the framework can be integrated into assignments, examinations, project-based learning, portfolios, and authentic assessments. Detailed case studies from different disciplines would provide practical examples of how educators can operationalise VOR and support student development of these capabilities.

Fourth, comparative studies should examine whether VOR-informed assessment approaches provide more valid evidence of learning than traditional product-centred assessment methods. Such research could investigate combinations of product, process, reflective, and oral evidence to determine which assessment designs best support learning, academic integrity, and responsible AI use.

Finally, future work should explore accessibility, fairness, inclusion, and ethics, including how technological competence, AI access, disciplinary cultures, linguistic backgrounds, and learner characteristics affect students' ability to demonstrate VOR capabilities. The framework's long-term value will depend not only on its theoretical coherence but also on its capacity to support equitable, inclusive, and educationally meaningful assessment in increasingly AI-mediated learning environments.

CONCLUSION

Generative AI does not make Bloom's Taxonomy outdated. It changes the conditions under which Bloom's cognitive processes occur and, crucially, the evidence educators use to infer learning. The six processes, Remember, Understand, Apply, Analyse, Evaluate, and Create, remain useful as a language for cognitive demand, but sophisticated output is no longer a sufficient proxy for sophisticated thinking when AI can participate directly in its production.

This paper therefore proposes a layered rather than replacement model. Bloom's six cognitive processes remain intact and are reinterpreted for AI-mediated learning. AI literacy functions as an enabling capability across the

model. Verify, Orchestrate, and Reflect operate as cross-cutting capabilities that make explicit the learner's relationship to information, technological resources, and their own reasoning. Their integration supports responsible Judgement, not as a seventh rung above Create, but as an emergent quality of higher-order performance.

The resulting shift is from creation as an unquestioned endpoint toward judgement surrounding creation. Creativity remains educationally important, but its value increasingly depends on whether learners can establish what to trust, decide how human and AI resources should contribute, explain and critique the process, and accept responsibility for the outcome. This reframing preserves Bloom's concern with cognitive complexity while addressing a central reality of generative AI: production can be shared, but educational responsibility cannot simply be delegated.

For assessment, the implication is equally clear. Universities should not respond to AI solely by trying to protect conventional products from technological assistance. They should design assessment systems that triangulate products with evidence of reasoning, verification, process and ownership. Viva examinations, authentic tasks, process documentation, AI critique and verification-focused activities are useful not because they are intrinsically AI-proof, but because they can reveal dimensions of learning that final artefacts alone may conceal.

The central question for higher education is therefore no longer simply, “Can the student produce an answer?” It is increasingly, “Can the student understand the problem, determine what to trust, use AI and other resources intelligently, explain the reasoning, and take responsibility for the outcome?” If Bloom is reinterpreted around that question, it can remain a valuable framework for the AI era while VOR provides a complementary vocabulary for the epistemic, strategic, and reflective capabilities that responsible human-AI learning now demands.

Disclosure statement: No potential conflict of interest was reported by the author.

Funding: The paper did not receive any external funding.

REFERENCES

1. Ayyoub, A.M., Khlaif, Z.N., Shamali, M., Abu Eideh, B., Assali, A., Hattab, M.K., Barham, K.A. and Bsharat, T.R.K. (2025) "Advancing higher education with GenAI: factors influencing educator AI literacy," Frontiers in Education (Lausanne), 10. Available at: https://doi.org/10.3389/feduc.2025.1530721.
2. Bacalso, F.D., Villamor, F.E., Paclipan, C.P., Orquia, C.J.B., Rendon, I.P., Albiso, A.T. and Himang, C.M. (2026) "Institutional approaches to generative AI management in higher education: a systematic review," Frontiers in Education (Lausanne), 11. Available at: https://doi.org/10.3389/feduc.2026.1814426.
3. Bloom, B. S., Engelhart, M. D., Furst, E. J., Hill, W. H., & Krathwohl, D. R. (1956). Taxonomy of educational objectives: Handbook I: Cognitive domain. David McKay.
4. Bobula, M. (2024) "Generative artificial intelligence (AI) in higher education: a comprehensive review of challenges, opportunities, and implications," Journal of Learning Development in Higher Education, 2024(30). Available at: https://doi.org/10.47408/jldhe.vi30.1137.
5. Hazari, S. (2024) "Justification and Roadmap for Artificial Intelligence (AI) Literacy Courses in Higher Education," Journal of Educational Research and Practice, 14(1). Available at: https://doi.org/10.5590/JERAP.2024.14.1.07.
6. Ilieva, G., Yankova, T., Ruseva, M. and Kabaivanov, S. (2025) "A Framework for Generative AI-Driven Assessment in Higher Education," Information (Basel), 16(6), p. 472. Available at: https://doi.org/10.3390/info16060472.
7. Jin, Y., Yan, L., Echeverria, V., Gašević, D. and Martinez-Maldonado, R. (2024) "Generative AI in Higher Education: A Global Perspective of Institutional Adoption Policies and Guidelines." Available at: https://doi.org/10.48550/arxiv.2405.11800.
8. Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet,

O., Sailer, M., Schmidt, A., Seidel, T., Stadler, M., Weller, J., Kuhn, J. and Kasneci, G. (2023) “ChatGPT for good? On opportunities and challenges of large language models for education," Learning and Individual Differences, 103, p. 102274. Available at: https://doi.org/10.1016/j.lindif.2023.102274.

9. Krathwohl, D.R. (2002) "A Revision of Bloom's Taxonomy: An Overview," Theory into Practice, 41(4), pp. 212–218. Available at: https://doi.org/10.1207/s15430421tip4104_2.
10. Lubbe, A., Marais, E. and Kruger, D. (2025) "Cultivating independent thinkers: The triad of artificial intelligence, Bloom's taxonomy and critical thinking in assessment pedagogy," Education and Information Technologies, 30(12), pp. 17589–17622. Available at: https://doi.org/10.1007/s10639-025-13476-x.
11. Miao, F. and Holmes, W. (2023) Guidance for generative AI in education and research. UNESCO (United Nations Educational, Scientific and Cultural Organization). London: UCL Knowledge Lab. Available at: https://discovery.ucl.ac.uk/id/eprint/10176438/1/Miao%20and%20Holmes%20-%202023%20-%20Guidance%20for%20Generative%20AI%20in%20Education%20and%20Resear.pdf (Accessed: 23 September 2026).
12. O'Dea, X. (2024) "Generative AI: is it a paradigm shift for higher education?," Studies in Higher Education (Dorchester-on-Thames), 49(5), pp. 811–816. Available at: https://doi.org/10.1080/03075079.2024.2332944.
13. Stanford Human-Centered Artificial Intelligence (2026) AI Index Report 2026: Education chapter. Stanford University. Available at: https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf (Accessed: 23 September 2026).
14. Stephenson, R. and Armstrong, C. (2026) Student Generative AI Survey 2026 (HEPI Report 199). Oxford: Higher Education Policy Institute (HEPI). Available at: https://www.hepi.ac.uk/wp-content/uploads/2026/03/HEPI-Report-199-Gen-AI-Survey-2026.pdf (Accessed: 23 September 2026).
15. Xia, Q., Weng, X., Ouyang, F., Lin, T.J. and 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(1), pp. 40–22. Available at: https://doi.org/10.1186/s41239-024-00468-z.