

Preserving Scientific Reasoning in Generative AI-Assisted Learning: A Proposed RPVR Framework for Basic Medical Sciences Education

Nurhaslina Hasan^{1*}, Vivi Noryati Ahmad¹, Faezah Sabirin¹, Zatifarihiyah Rasdi¹

Centre of Preclinical Science Studies, Faculty of Dentistry, Universiti Teknologi Mara, Sg Buloh
Campus, 47000 Sg Buloh, Selangor, Malaysia.

*Corresponding Author

DOI: <https://doi.org/10.47772/IJRISS.2026.102100184>

Received: 20 September 2026; Accepted: 25 September 2026; Published: 07 October 2026

ABSTRACT

Integration of generative artificial intelligence (GenAI) in higher education offers substantial opportunities for personalised education that tailors to the needs, interests and learning of each student. However, unrestricted reliance on GenAI has been found to be a source of academic stress and is associated with low self-esteem. Moreover, in subjects that require students to develop critical thinking or scientific reasoning, such as basic medical sciences (BMS), high dependence on AI may limit students' opportunities to engage in independent knowledge construction and critical evaluation. This concept paper proposes an integration of GenAI into BMS education through the Reason–Prompt–Verify–Reconstruct (RPVR) framework. The framework's approach is at both the learner and educator levels, with the main aims of preserving active scientific reasoning. Comprises of four sequential stages: R Reason, stage where learners independently analyse a problem and formulate an initial explanation before using GenAI; P Prompt, a stage where the use of GenAI is allow to challenge, clarify, or extend the learner's reasoning; V Verify, stage which requires critical appraisal of AI-generated information against standard academic and scientific sources; and R Reconstruct, where learners identify misconceptions, reflect their understanding independently, and transfer the verified concept to a novel problem. Problem-based learning (PBL) is proposed to illustrate the implementation of RPVR, using a sequence of AI-off–AI-on–verification–AI-off to ensure that learners generate initial hypotheses before AI integration and subsequently demonstrate understanding through explanation, mechanistic reconstruction, and transfer to unfamiliar scenarios. As physiology strongly depends on causal and mechanistic reasoning, this subject is used as the primary illustrative discipline, while the framework is potentially transferable to any other BMS disciplines. By repositioning GenAI from an answer-generating tool, RPVR aims to promote adaptive cognitive offloading and broaden the use of AI in higher education. Future empirical studies should evaluate whether RPVR-guided GenAI use, particularly within PBL, improves reasoning, knowledge transfer, and learning outcomes compared with unstructured GenAI-assisted learning.

Keywords: AI-assisted learning; basic medical sciences education; problem-based learning; physiology education.

INTRODUCTION

Generative artificial intelligence (GenAI) is increasingly used in higher education for tutoring, personalised explanations, feedback, self-assessment, and content support. Reviews indicate rapid expansion of these adoptions, but evidence for GenAI-specific pedagogical approaches remains immature (Basil et al., 2026; Simoni et al., 2025). While GenAI is increasingly adopted in medical education, its systematic deep integration into basic medical sciences (BMS) remains limited. Generally, BMS encompasses disciplines such as physiology, biochemistry, pharmacology, and pathology, which cover the conceptual and mechanistic foundations for understanding normal biological function to disease processes, implications, and management. Students must progressively develop the ability to identify causal relationships, construct mechanistic explanations, integrate concepts across systems, and apply foundational knowledge to interpret clinical

consequences. Effective learning in the BMS, therefore, requires more than memorising facts. Recent work in pharmacology education illustrates this distinction clearly. Bloch et al. (2025) demonstrated that medical students who constructed more complete causal-mechanistic explanations were substantially more likely to correctly predict an adverse drug effect. The finding supports the educational importance of reasoning through underlying mechanisms rather than relying solely on recall or recognition.

GenAI can facilitate such learning by clarifying difficult concepts, generating examples, and responding interactively to student questions. A customised GPT physiology companion, for example, has demonstrated the feasibility of embedding AI tutoring alongside conventional teaching, with students reporting favourable perceptions of its educational values (Elsamanoudy et al., 2026). Such capabilities may enhance access to learning support, particularly when students encounter complex facts or in need additional clarification. Nonetheless, the students may also have access to a fully developed answer before they attempt the given task independently. Nevertheless, current evidence does not establish that GenAI inherently enhances or undermines higher-order reasoning. Systematic and scoping reviews highlighted the limited theory-informed research and a lack of direct evidence on critical thinking or clinical reasoning outcomes from GenAI adoption (Feigerlova et al., 2025; Simoni et al., 2025). The central concern is therefore not that GenAI necessarily diminishes learning. Instead, the concern is that the unstructured adoption of GenAI may affect the cognitive process by limiting students' opportunities to generate, test, and refine their own scientific reasoning. In hand, GenAI creates both an opportunity and a pedagogical dilemma within this context. The pedagogical concern is therefore premature or unstructured reliance rather than the use of GenAI itself.

Problem-based learning (PBL) is a pedagogical approach that aims to provide a relevant context for educational value, relying on learners' ability to encounter unfamiliar problems, activate prior knowledge, generate hypotheses, identify knowledge gaps, and reconstruct explanations. Recent research on physiology education suggests that well-structured PBL can strengthen causal and mechanistic reasoning even before formal teaching of the relevant topics (Morales-Luque et al., 2026). A 2025 systematic review and meta-analysis reported an overall advantage of PBL over conventional teaching for critical-thinking outcomes, while also noting substantial heterogeneity associated with educational context and implementation (Su et al., 2025). Emerging studies have begun integrating ChatGPT into PBL and have reported promising improvements in learning outcomes, satisfaction, and efficiency in selected clinical settings (Tong et al., 2025). Yet, immediate access to a complete AI-generated synthesis, especially among novice learners, may bypass the very processes that PBL is intended to develop. The pedagogical challenge is therefore to preserve the productive uncertainty and independent reasoning that underpin PBL while still allowing students to develop the ability to work responsibly with GenAI.

Several frameworks have already been proposed to provide important guidance on integrating GenAI into education (Mengesha et al., 2026; Su & Yang 2023). Nevertheless, a specific need remains for a learner-level framework that protects independent scientific reasoning before AI engagement and requires verification of AI-generated scientific information before acceptance. Such a framework should also ensure that students independently reconstruct and transfer their understanding after AI adoption. Thus, this paper addresses how GenAI can be integrated into BMS while preserving independent scientific reasoning and critical evaluation. The Reason–Prompt–Verify–Reconstruct (RPVR) framework is proposed with PBL as a principal implementation environment and physiology as the primary illustrative discipline.

LITERATURE REVIEW

In general, basic medical sciences (BMS) form an essential foundation for subsequent clinical application in healthcare education. Their educational value lies not only in the quantity of information retained but also in the student's ability to integrate the information through defensible explanatory chains. Factual correctness is insufficient when students are unable to explain why an outcome occurs or predict how the outcome changes under a novel condition. In physiology, students are required to integrate multiple variables and understand the underlying mechanisms, for example, how exercise shifts the oxyhaemoglobin dissociation curve to the right. The knowledge is not just about identifying the factors affecting and the consequences. Instead, a deeper understanding requires the ability to explain how exercise affects factors such as temperature, pH and carbon

dioxide, which subsequently alter haemoglobin–oxygen affinity, thereby regulating oxygen supply. Similarly, in pharmacology, causal-mechanistic reasoning has been found to support prediction rather than recognition alone (Bloch et al., 2025).

GenAI has the characteristics that can support BMS educational aims, which provide the constructional, refinement, and transfer of explanations rather than simply increasing access to the information. GenAI, through its conversational model systems, can adjust explanatory depth, generate examples, pose questions and provide rapid formative support, outside formal teaching sessions. These capabilities support personalised and self-directed learning, unlike static educational resources such as textbooks or videos that have fixed content. A recent review by Simoni et al. (2025) has documented a wide range of AI applications in undergraduate medical education (UGME), from tutoring, self-knowledge assessment and skill assessment. Nonetheless, none of the reviewed articles fosters AI in higher-order reasoning in UGME. A similar pattern was observed in AI integration in other health professions education, which reported rapid adoption but insufficient evidence on complex educational outcomes such as critical thinking (Basil et al., 2026; Feigerlova et al., 2025).

The educational benefits of GenAI come along with several potential cognitive and epistemic risks. One concern is cognitive offloading, whereby learners delegate components of a task to an external tool rather than performing those cognitive processes independently. Cognitive offloading is not inherently undesirable, as certain tools may reduce unnecessary cognitive burden. The issue arises when the cognitive process is totally outsourced, which is the capability that students are expected to develop. The second concern involves the reliability of AI-generated information. Large language models generate plausible linguistic responses but may produce incomplete, inconsistent, and, in the worst case, inaccurate information. Feigerlova et al. (2025) noted variable accuracy across AI educational models and emphasised the need for educators to consider inaccuracies when designing AI-supported strategies. Despite focusing on the favourable results from ChatGPT-guided PBL, Tong et al. (2025) also acknowledged the potential for misleading or incorrect information and the importance of appropriate prompting. Consequently, students require more than technical proficiency in prompting, where they must also evaluate the epistemic quality of AI-generated scientific claims. Moreover, GenAI responses are often presented fluently and confidently, thus, novice learners may interpret consistency of presentation as evidence of scientific accuracy. The educational response should therefore not rely on simply telling students to check ChatGPT, but on making verification an explicit learning behaviour in which claims are compared with reputable materials such as established textbooks, lecturer-provided materials, or other disciplinary sources.

It is important that the aforementioned concerns should not be interpreted as evidence that GenAI inherently damages critical thinking. Current empirical evidence remains insufficient for such a conclusion. Simoni et al. (2025) reported that, despite rapid growth of AI use across undergraduate medical education, none of the publications included in their updated scoping review directly evaluated AI's impact on critical thinking or clinical reasoning. The appropriate conclusion is therefore more cautious: unstructured or premature dependence on GenAI may reduce opportunities for independent reasoning, whereas appropriately structured GenAI interaction may scaffold that reasoning. Therefore, the educational challenge is to design a learning sequence that distinguishes between these two possibilities, allowing learners to utilise GenAI while preserving their independent scientific reasoning (Pham et al., 2025).

Problem-based learning (PBL) provides a particularly relevant environment for addressing this challenge. Unlike conventional teaching approaches, PBL begins with a problem that learners are expected to analyse using their existing knowledge. Through discussion, hypothesis generation, identification of knowledge gaps, self-directed inquiry, and subsequent synthesis, students progressively construct an explanation of the problem. This structure makes students' reasoning processes central to learning rather than treating them solely as a pathway towards obtaining the correct final answer. Evidence generally supports PBL's capacity to promote higher-order thinking, although the outcomes depend on its conduct. Su et al. (2025) appraised 11 studies and reported an overall advantage of PBL over conventional teaching for critical-thinking outcomes, while noting substantial heterogeneity related to educational context, assessment methods, and implementation fidelity.

GenAI can potentially strengthen several components of PBL. It may help learners explore alternative explanations, generate questions, clarify unfamiliar concepts, compare hypotheses, or obtain rapid formative feedback. Emerging empirical studies support this possibility. Tong et al. (2025), for example, compared ChatGPT-guided PBL with traditional teaching among nephrology trainees and reported higher theoretical and clinical assessment scores, together with greater satisfaction, in the intervention group. These findings provide early evidence that GenAI and PBL may be pedagogically compatible.

However, integrating GenAI into PBL without regulating its timing may alter the mechanism through which PBL is expected to produce learning. PBL deliberately places students in a state of incomplete knowledge that requires them to formulate and revise explanations. If GenAI resolves that uncertainty immediately after the case is presented, learners may bypass hypothesis generation and mechanistic construction. The risk is therefore not that AI provides information, since information seeking is already integral to PBL, but that AI can provide a complete synthesis before students have articulated what they currently think, what they do not understand, and why their initial explanation may be incomplete. This suggests a need for structured phases of AI access. Students should first analyse the issue independently, commit to an initial explanation, and make their reasoning visible prior to using GenAI. GenAI could subsequently be used to challenge rather than replace the explanation. The GenAI output should then be verified against established evidence before AI access is withdrawn. Lastly, students are required to independently reconstruct and transfer their understanding through either a written or a verbal assessment. This sequence provides the conceptual bridge between established PBL principles and the proposed RPVR framework.

Several frameworks provide useful foundations for responsible GenAI integration. Su and Yang (2023) IDEE framework was among the early theoretical models developed specifically for the educational use of GenAI. IDEE involves identifying desired outcomes, determining the appropriate level of automation, ensuring ethical considerations, and evaluating effectiveness. Its major contribution is to position AI adoption as an educational-design decision rather than a technology-driven decision, emphasising that learning objectives should precede tool selection. On the other hand, Acar's PAIR framework provides a more learner-facing approach comprising Problem, AI, Interaction, and Reflection. Students first formulate a problem, select or engage with an appropriate AI tool, interact iteratively with the tool, and reflect on the experience (Acar, 2023). The model emphasises human-centred, skill-centred, and responsible AI use and has subsequently been piloted and adapted in higher-education settings. PAIR is especially relevant to the present concept because it recognises that educational AI use should involve active human judgement rather than passive acceptance of generated outputs. The recently proposed ASPIRE framework provides a more comprehensive curriculum-level model (Mengesha et al., 2026). Developed within public-health education, ASPIRE comprises six interconnected phases: Assess, Scaffold, Plan, Implement, Reflect, and Evaluate. It integrates constructive alignment, scaffolded learning, AI literacy, assessment design, ethical judgement and reflection, and explicitly encourages educators to design assessment tasks that make students' thinking visible. This represents an important development from general guidance towards theoretically grounded and structurally embedded GenAI pedagogy.

Overall, the existing GenAI frameworks provide important guidance regarding educational alignment, human–AI interaction, ethical use, reflection, and assessment design. The distinction is important as PAIR begins with problem formulation before proceeding to AI engagement, while RPVR proposes that learners should additionally construct and make visible an initial disciplinary explanation before AI exposure. Furthermore, whereas critical evaluation may occur within interaction in existing models, RPVR positions verification as a distinct mandatory phase in which AI-generated scientific claims must be compared against authoritative disciplinary evidence before they are incorporated into the learner's understanding. Finally, reconstruction within RPVR is intended not merely as reflection on the AI experience but as an independent reconstruction and transfer of the scientific concept. This conceptual foundation provides the basis for the subsequent discussion, in which the RPVR framework, its application within PBL, and a physiology-based worked example are presented in detail.

DISCUSSION

The Proposed Reason–Prompt–Verify–Reconstruct (RPVR) Framework

The increasing availability of GenAI in providing immediate explanations, examples, feedback, and alternative perspectives creates an important pedagogical tension in higher education, including BMS. The accessibility may allow learners to obtain a complete explanation before engaging sufficiently with the reasoning processes required to construct it. The educational challenge is therefore not simply whether students should use GenAI, but how its use can be sequenced so that technology augments rather than substitutes for scientific reasoning. To address this challenge, this concept paper proposes the Reason–Prompt–Verify–Reconstruct (RPVR) framework, a four-stage learner-centred approach that regulates the timing and cognitive purpose of GenAI interaction.

RPVR focuses specifically on preserving observable disciplinary reasoning within an AI-supported learning sequence. The framework is organised around four sequential stages: Reason, Prompt, Verify, and Reconstruct. The sequence deliberately begins and ends with predominantly human cognitive activity. GenAI is introduced only after learners have made their initial reasoning visible and is subsequently withdrawn after its output has been critically verified. Thus, the proposed learning sequence may be represented as:

Human reasoning → AI-supported interrogation → evidence verification → human reconstruction and transfer rather than:

Problem → AI-generated answer → learner reproduction

This sequencing is particularly relevant to BMS, where the educational objective is not merely to identify a precise fact but to explain the causal and mechanistic relationships through which that outcome occurs.

Stage 1: Reason — Making Independent Thinking Visible

The first stage of RPVR requires learners to analyse a problem and construct an initial explanation before accessing GenAI. The instructor establishes the AI-off condition, clarifies the task, and may use non-directive probing questions without supplying the explanation. The purpose of this stage is not to ensure that the learner's first explanation is precise or complete. Instead, it is intended to activate prior knowledge, reveal assumptions, identify conceptual gaps, and make the learner's reasoning available for subsequent evaluation. During this stage, students may be required to identify relevant variables, predict the direction of change, propose causal relationships, formulate hypotheses, or construct a simple mechanism map. Importantly, students should also indicate areas of uncertainty. A brief structured response may therefore include:

- the learner's initial prediction;
- the proposed mechanism;
- evidence or prior knowledge supporting the prediction; and
- concepts about which the learner remains uncertain.

For example, when presented with a patient experiencing metabolic acidosis, students might initially be asked:

“Predict the effect on the oxyhaemoglobin dissociation curve and explain the physiological mechanism underlying this change.”

Rather than immediately consulting GenAI, learners would first construct a pathway such as:

$\uparrow \text{H}^+$ ($\downarrow \text{pH}$, acidosis) → altered haemoglobin–oxygen interaction (haemoglobin as acid-base regulator) → $\downarrow$ haemoglobin affinity for O_2 → rightward shift → increased oxygen unloading to metabolically active tissues

The subsequent stage can therefore interact with this existing cognitive representation rather than replacing the need to create one. This stage also has an important assessment advantage as capturing the initial explanation before AI exposure creates a visible baseline against which later conceptual development can be evaluated.

Stage 2: Prompt — Using GenAI as a Cognitive Scaffold

Once learners have articulated an initial explanation, GenAI becomes available during the Prompt stage. However, its pedagogical function should be explicitly defined. Instead of functioning primarily as an answer generator, GenAI should serve as a cognitive scaffold that interrogates challenges, extends, or tests the learner's reasoning. This includes identifying a missing relationship, challenging an assumption, requesting an alternative explanation, or generating a counterexample rather than asking for a direct complete answer. Prompt quality therefore becomes part of responsible AI literacy. The instructor may model appropriate prompting and redirect interaction when learners begin using GenAI primarily for answer retrieval. For example, instead of asking;

“Explain the effect of acidosis on the oxyhaemoglobin dissociation curve.”

learners may be encouraged to prompt:

“This is our explanation of how acidosis influences haemoglobin–oxygen affinity. Identify weaknesses or missing steps without giving us a complete replacement answer.”

Other useful prompt structures may include:

- “Ask three questions that would test whether the provided mechanism is scientifically precise”
- “Identify a variable that our explanation has overlooked.”
- “Challenge one assumption in our explanation.”
- “Give us a counterexample that would require us to modify our reasoning.”

Such prompts reposition GenAI from an authoritative knowledge source to an interactive reasoning partner. This is consistent with the concept of scaffolding which support is provided in a manner that enables learners to progress beyond their initial understanding without removing responsibility for constructing that understanding. The quality of prompting itself may therefore become an educational competency. Students should learn that effective GenAI use in science is not equivalent to obtaining the most complete answer with the fewest words. Instead, sophisticated AI use may involve generating prompts that expose uncertainty, elicit alternative explanations, identify assumptions, and stimulate deeper disciplinary reasoning.

Stage 3: Verify — Treating AI Output as a Claim Rather Than Evidence

The Verify stage represents a central feature of RPVR. GenAI-generated statements should not automatically be treated as established scientific evidence, regardless of how fluent or convincing the response appears. Students must therefore critically compare important AI-generated claims with authoritative disciplinary sources before incorporating them into their final understanding. Verification may involve prescribed textbooks, lecturer-provided learning materials, peer-reviewed literature, recognised professional guidelines, or validated scientific databases. The specific sources should be appropriate to the level of the learner and the complexity of the task. A simple verification matrix could be incorporated into learning activities, for example, as shown in Table 1.

Table 1 Verification matrix of AI-generated claim

AI-generated output	Establish source	Students' judgement
Acidosis decreases haemoglobin affinity for oxygen	Prescribed physiology textbook	Supported
The Bohr effect primarily increases oxygen loading	Textbook/lecture materials	Incorrect; modification

in tissues		required
Increased tissue CO ₂ contributes to greater oxygen unloading	Textbook/peer-reviewed source	Supported

The educational purpose is not merely to detect AI “hallucinations”. Verification develops a broader form of critical AI literacy, in which students learn to distinguish between plausible language and defensible scientific knowledge. Learners must justify why an AI-generated statement is accepted, modified, or rejected. This distinction is particularly important for novice students. Experts can often identify subtle inaccuracies because they already possess extensive disciplinary knowledge. Beginning learners, however, may be more vulnerable to accepting confident but incomplete explanations. Consequently, verification should not be treated as optional advice; within RPVR, it becomes an explicit and observable component of the learning process.

Stage 4: Reconstruct — Reflecting and Transferring Understanding

The final stage of RPVR is Reconstruct, which is defined more comprehensively than simply considering whether the AI response was useful. Within RPVR, reflection involves three linked activities:

Reconstruct → Reflect → Transfer

First, students compare their initial explanation with the subsequently verified understanding. They identify misconceptions, missing mechanisms, unsupported assumptions, and changes in their thinking. This encourages metacognitive awareness of how their understanding developed. Second, learners should reconstruct the explanation independently, ideally without access to GenAI. This requirement is crucial because a student who can recognise a correct AI-generated explanation may not necessarily be able to reproduce the underlying reasoning independently. Third, students should demonstrate transfer by applying the same underlying principle to a different problem. Transfer provides stronger evidence of conceptual understanding than repeating the original explanation. For example, after investigating the effect of acidosis on oxygen–haemoglobin affinity, students may be given the unfamiliar question:

“A patient undergoing prolonged surgery develops significant hypothermia. Predict the resulting change in haemoglobin–oxygen affinity and explain its potential effect on tissue oxygen unloading.”

Learners must now apply their understanding of factors affecting the oxyhaemoglobin dissociation curve rather than retrieve the answer previously generated by AI. A learner who merely memorised the first explanation may struggle, whereas a learner who understands the underlying mechanism should be better able to reason through the new scenario. Accordingly, the Reconstruct stage transforms RPVR from an information-verification model into a learning and transfer framework.

Implementation

Problem-based learning (PBL) provides a particularly appropriate context for implementing RPVR because its educational value depends on learners actively engaging with uncertainty, generating hypotheses, identifying knowledge gaps, seeking relevant information, and reconstructing explanations collaboratively. RPVR can regulate GenAI involvement without fundamentally changing the established principles of PBL. The proposed implementation can be conceptualised as an:

AI-off → AI-on → verification → AI-off sequence

Following the presentation of the PBL trigger, students initially work without GenAI. Each learner may first complete a brief individual reasoning task before group discussion. Before the discussion and AI consultation, each student should record a brief individual explanation. This creates individual cognitive accountability and reduces the likelihood of passive participation, or the stronger group members dominating the initial interpretation. The group then:

- identifies relevant facts;
- generates hypotheses;
- constructs an initial mechanism;
- identifies knowledge gaps; and
- formulates learning issues.

Only after the group has documented its preliminary reasoning does GenAI become available. During the second phase, students use GenAI to challenge rather than replace their reasoning. Access may initially be managed through a designated AI navigator who enters group-approved prompts. Limiting AI interaction to one visible interface can make the process transparent and encourage discussion about what should be asked and why. Appropriate GenAI functions may include:

- identifying missing mechanisms;
- questioning assumptions;
- generating alternative hypotheses;
- producing counterexamples;
- asking Socratic questions; and
- suggesting concepts requiring further investigation.

GenAI-generated text should not constitute the group's final response. The group then evaluates relevant AI-generated claims using authoritative sources. Statements may be classified as either supported, partially supported, inaccurate, incomplete, or uncertain. This stage transforms PBL information seeking from passive retrieval into evidence appraisal. GenAI is subsequently withdrawn. Students must reconstruct their final explanation in their own words and demonstrate their understanding through discussion, mechanism mapping, tutor questioning, or an unfamiliar transfer problem.

This final AI-free phase is particularly important because it provides evidence that learning has been internalised rather than temporarily supported by access to GenAI. Thus, PBL provides the learning environment, while RPVR regulates when and how GenAI enters the reasoning process. A major practical concern in AI-supported PBL is that students may simply read or paraphrase a generated answer during reporting. RPVR, therefore, requires instructional design strategies that make independent understanding necessary rather than optional. Students should not be permitted to present an AI-generated passage as the group's explanation. The final response must be reconstructed in the group's own words and preferably organised around causal or mechanistic relationships rather than narrative text.

During reporting, facilitators may randomly select any group member to explain a particular component of the mechanism. Students, therefore, cannot rely solely on a designated presenter. Each PBL session should conclude with a new question that cannot be answered merely by reading the previously generated AI response. Transfer questions can involve changing one or more physiological variables, introducing a new disease state, or asking students to predict the consequences of failure of a compensatory mechanism. Transfer is therefore positioned as a proof of understanding, rather than an optional extension activity.

Although physiology is used as the primary example, the RPVR sequence is potentially applicable across BMS disciplines because many learning tasks involve explanation, prediction, evidence appraisal, and transfer. The cognitive emphasis may differ between disciplines. Physiology and pharmacology may depend strongly on causal-mechanistic reasoning, anatomy may emphasise spatial and structural reasoning, and microbiology may require integration of organism, host, and antimicrobial factors. Nevertheless, the broad sequence remains applicable:

attempt reasoning → use AI strategically → verify claims → demonstrate independent understanding

Accordingly, RPVR should be considered a flexible framework rather than a rigid scripted protocol.

Assessment

If educators value independent reasoning but assessments reward only the final correct answer, students may have little incentive to engage deeply with the RPVR process. Assessment should therefore make at least part of the reasoning process visible. Table 2 shows an example of a formative RPVR rubric with these allocated percentages are illustrative rather than prescriptive. The key principle is that AI use itself should not be rewarded merely because students used a sophisticated tool. Credit should instead reflect how effectively learners retain intellectual responsibility for the problem.

Table 2 Formative assessment rubric of RPVR implementation in PBL

Domain	Illustrative weighting
Independent initial reasoning	25%
Quality and coherence of mechanistic explanation	25%
Purposeful use of GenAI prompting	15%
Verification and critical appraisal	15%
Reconstruction and transfer to a novel problem	20%

Assessment evidence might include:

- initial reasoning maps;
- AI prompt logs;
- verification tables;
- corrected mechanism maps;
- brief reconstructive statements;
- oral teach-back; and
- performance on unseen transfer questions.

This process-based approach may also provide greater transparency regarding how GenAI contributed to an assignment or learning activity.

Role of the Educator

RPVR does not diminish the importance of the educator, rather it defines it. The adoption of GenAI changes the educator's role from predominantly providing information to designing cognitive conditions in which students must reason, challenge, verify, and transfer the knowledge. As GenAI increasingly provides access to information, explanations, and feedback, the educator's primary responsibility shifts from being a major source of content delivery to designing and facilitating learning environments that promote reasoning, critical evaluation, and knowledge transfer. This perspective aligns with contemporary views of GenAI as a pedagogical tool whose educational value depends on thoughtful instructional design rather than the technology itself (Su & Yang 2023). Facilitators are responsible for:

- selecting problems that require explanation rather than simple recall;
- determining when AI access is appropriate;
- modelling effective scientific prompting;
- identifying authoritative verification sources;
- challenging overconfident AI-derived statements;
- facilitating discussion of uncertainty;
- ensuring equitable group participation; and
- designing transfer questions that reveal genuine understanding.

Educators also require sufficient AI literacy to recognise both useful and inappropriate applications of GenAI. RPVR therefore implies parallel development of learner and educator competencies.

Distinction of the RPVR Framework from the Existing Framework

RPVR is not intended to replace existing GenAI educational frameworks. The framework addresses a more specific learner-level guideline for solving scientific problems that preserves students' cognitive abilities and accomplishes higher-order thinking before, during, and after interacting with GenAI. The comparison between the proposed and existing frameworks is summarised in Table 3.

Table 3 Comparison of the related GenAI educational framework with the RPVR framework

Framework / Feature	IDEE (Su & Yang, 2023)	PAIR (Acar, 2023)	ASPIRE (Mengesha et al., 2026)	UGME Curriculum	Proposed RPVR
Focus	Responsible integration of GenAI	Structured learner interaction with GenAI	Curriculum/ and assessment design for GenAI integration	UGME curriculum - integration in competencies	Preservation of learner higher-order thinking in GenAI integration into BMS
Level of application	Educator – instructional design	Learner – AI interaction	Educator – curriculum and assessment design	Curriculum/ programme level	Learner – AI interaction focuses on cognitive
Core structure	Identify outcomes → Determine automation → Ensure ethics → Evaluate effectiveness	Problem → AI → Interaction → Reflection	Assess → Scaffold → Plan → Implement	Seven interconnecte d domains, including AI knowledge, application, human-AI collaboration, ethics, research collaboration and assessment	Reason (scientifically) → Prompt (AI) → Verify → Reconstruct/ Reflect
Point of AI integration	Determined according to instructional goals	Following problem formulation and AI tools selection	Determined during curriculum and activity design	Not prescribed as a fixed learner sequence	After learners' initial scientific explanation
Learner-generated reasoning before AI	Not explicitly required	Problem formulation is required, but an independent explanation before AI use	Student reasoning may be designed into learning and assessment, but no fixed	Higher-order thinking is incorporated through curriculum domains, but no mandatory	Explicit requirement: learners construct and document their initial reasoning before AI exposure

		is not explicitly mandated	pre-AI reasoning stage is prescribed	pre-AI reasoning stage is specified	
GenAI interaction	Depends on the selected educational application	Learners experiment with AI inputs and outputs	GenAI interaction is scaffolded according to learning and assessment design	Human-AI collaboration is one curriculum domain	AI is deliberately prompted to challenge, question, clarify or extend learner-generated reasoning rather than immediately replace it
Critical evaluation of AI output	Included within ethical/effectiveness considerations	Explicitly emphasised; learners are advised to critically evaluate and corroborate outputs	Strong emphasis on evidence, judgement, assessment validity and responsible use	Evaluation, assessment and human-AI collaboration are included	Explicit and compulsory before AI information is incorporated into the learner's final explanation
Verification against authoritative disciplinary evidence	Not prescribed as a fixed learner stage	Recommended through corroboration with trusted sources	May be incorporated through evidence-based learning and assessment	Relevant to evaluation and assessment but not defined as a sequential learner stage	A distinct stage requiring accept-modify-reject decisions based on authoritative disciplinary evidence
AI withdrawal	Not specified	Not explicitly required	Not explicitly required	Not specified	Explicit: AI is withdrawn before the learner demonstrates final understanding
Post-AI independent reconstruction	Not specified	Reflection focuses substantially on the AI interaction and problem-solving experience	Student reflection can be incorporated within assessment; educator reflection is also central	No prescribed post-AI learner sequence	Required: learners reconstruct the scientific explanation independently after AI withdrawal
Transfer to a novel problem	Not explicitly prescribed	May occur depending on activity design, but is not a required framework stage	Can be incorporated into assessment design but is not intrinsic to the sequence	Higher-order application is compatible with the framework but not prescribed as a post-AI stage	Explicitly used to demonstrate that understanding extends beyond the AI-supported task
Distinguishes	Not a central	Not explicitly	Assessment	Evaluation	Central feature:

AI-supported performance from independent learning	focus	operationalised through AI withdrawal	validity and evidence of learning are major concerns, but no fixed withdrawal sequence is prescribed	and assessment are curriculum domains	performance during AI access is distinguished from understanding demonstrated after AI withdrawal
Assessment role	Evaluate whether GenAI achieves intended outcomes	Reflection/reporting of AI use can be incorporated into coursework	Major focus on assessment redesign, validity and making reasoning visible	Evaluation and assessment are integral curriculum domains	Captures a before-during-after reasoning trace and can align independent reconstruction/transfer with assessment
Educator role	Select appropriate goals, automation and ethical safeguards	Guide appropriate and responsible learner use of AI	Central role in scaffolding, curriculum design, assessment and evaluation	Design and deliver AI-integrated medical curricula	Design reasoning problems, control AI-entry points, guide verification, and assess independent reconstruction/ transfer
Disciplinary specificity	General education	Cross-disciplinary	Developed in public health education	Undergraduate medical education broadly	Designed for scientific and mechanistic reasoning, particularly in BMS
Principal distinctive contribution	Aligning GenAI use with outcomes, automation, ethics and evaluation	Developing responsible human-AI problem-solving skills	Systematically embedding GenAI across the curriculum and assessment	Establishing domains needed for AI-integrated undergraduate medical curricula	Cognitive gating of AI around two protected phases of human reasoning: learner-generated explanation before AI, and independent reconstruction/ transfer after AI withdrawal

It is important to highlight that RPVR does not claim that reasoning, verification, reconstruction or responsible AI use are novel educational principles. The proposed contribution lies in its deliberate sequencing around protected pre- and post-AI phases of learner-generated scientific reasoning.

LIMITATIONS OF THE PROPOSED FRAMEWORK

Several limitations should be acknowledged. First, RPVR is presently a conceptual framework and has not yet been empirically validated. Its effectiveness in improving scientific reasoning, retention, transfer, or AI literacy cannot therefore be assumed. Second, delaying AI access may not be necessary for every learning task. If the intended outcome is straightforward information retrieval or terminology clarification, requiring a complete RPVR cycle may create unnecessary cognitive and instructional burden. The framework is likely to be most relevant when the intended learning outcome involves reasoning, explanation, evidence evaluation, or transfer.

Third, implementation may be affected by differences in students' prior knowledge, AI literacy, access to technology, language proficiency, and familiarity with authoritative sources. Novice learners may require

greater scaffolding during verification than advanced students. Fourth, GenAI technologies evolve rapidly. The accuracy and functionality of future systems may improve substantially; however, improved accuracy would not eliminate the pedagogical need for students to retain responsibility for scientific reasoning and evidence appraisal. Last but not least RPVR is proposed across BMS but is illustrated principally using physiology. Applications in anatomy, biochemistry, pathology, microbiology, and pharmacology should therefore be considered potentially transferable rather than empirically established.

Future studies should compare RPVR-guided GenAI-assisted learning with unstructured GenAI use and, where appropriate, conventional learning without GenAI. Outcomes should extend beyond immediate factual performance to causal-mechanistic reasoning, verification accuracy, transfer to unfamiliar problems, delayed retention, critical AI literacy, metacognitive awareness, cognitive load, engagement, and learner perceptions. Within PBL, discussion quality, hypothesis generation, participation, prompt quality, verification behaviour, and the nature of lecturer facilitation are also relevant. Implementation fidelity should be examined to determine whether lecturer guidance supports verification and reconstruction without replacing learner reasoning. Crucially, researchers should distinguish performance while AI is available from independent performance after AI withdrawal; durable learning is more convincingly demonstrated when students can reconstruct and transfer knowledge without continued AI support.

CONCLUSION

GenAI is becoming increasingly embedded in the learning environment of students in the BMS. Its educational value should therefore not be considered solely in terms of the accuracy or convenience of the answers it generates. A more important question is whether its use preserves the cognitive processes through which learners develop scientific understanding. This concept paper proposes the Reason–Prompt–Verify–Reconstruct (RPVR) framework as a structured approach to responsible GenAI-assisted learning. RPVR positions independent reasoning before AI engagement, uses GenAI as a scaffold to challenge and extend learner-generated explanations, requires verification of AI-generated scientific claims against authoritative evidence, and concludes with reconstruction, reflective and transfer of understanding. Within PBL, this can be operationalised through an AI-off–AI-on–verification–AI-off sequence, thereby preserving the productive uncertainty and active reasoning central to problem-based learning while enabling students to develop responsible GenAI literacy.

Physiology provides a useful illustrative context because of its strong reliance on causal and mechanistic reasoning, although the framework is potentially transferable across other BMS disciplines. RPVR is not proposed as an alternative to existing educational AI frameworks, but as a complementary learner-level approach addressing a specific pedagogical question: when should GenAI enter the scientific reasoning process, and what should learners be required to do before and after that interaction. Ultimately, the educational objective should not be to prevent students from obtaining answers from GenAI, but to ensure that access to those answers does not replace the development of understanding. RPVR therefore proposes a shift from AI as an answer provider to AI as a reasoning scaffold, while retaining human reasoning, scientific verification, and knowledge transfer at the centre of learning. Empirical validation is now required to determine whether this structured approach improves scientific reasoning, critical AI literacy, and durable learning outcomes across BMS education.

REFERENCES

1. Acar, O. A. (2023). Are Your Students Ready for AI? A Four-Step Framework to Prepare Learners for a ChatGPT World. Harvard Business Publishing Education. <https://hbsp.harvard.edu/inspiring-minds/are-your-students-ready-for-ai?>
2. Basil, M., et al. (2026). A scoping review of the use of generative artificial intelligence tools in health profession education. BMC Medical Education, 26(1), 291.
3. Bloch, R. R., et al. (2025). Causal Mechanistic Reasoning as a Tool to Explore Medical Students' Predictions of Pharmacology Phenomenon: Connecting Core Concepts with Clinical Applications. Medical Science Educator, 35(5), 2329-2342.

4. Cheng, D., Chan, E. Y. S., Song, Y., Cui, A., Chang, A., & Wang, W. (2026). Framework for AI integration into the undergraduate medical curricula: a scoping review. *BMC medical education*, 26(1), 465. <https://doi.org/10.1186/s12909-026-08620-1>
5. Elsamanoudy, A., et al. (2026). Implementation and evaluation of customized artificial intelligence personal tutor (chat GPT physiology companion) for medical students in the reproduction module. *BMC Medical Education*, 26(1), 832.
6. Feigerlova, E., et al. (2025). A systematic review of the impact of artificial intelligence on educational outcomes in health professions education. *BMC Medical Education*, 25(1), 129.
7. Mengesha, Z., et al. (2026). ASPIRE: a pedagogical framework for integrating generative artificial intelligence in public health education. Paper presented at the *Frontiers in Education*.
8. Morales-Luque, C., et al. (2026). Problem-based learning promotes causal physiological reasoning in first-year medical students before formal cardiovascular instruction. *Frontiers in Physiology*, 17, 1897473.
9. Pham, T. D., et al. (2025). The impact of generative AI on health professional education: A systematic review in the context of student learning. *Medical Education*, 59(12), 1280-1289. doi:<https://doi.org/10.1111/medu.15746>
10. Simoni, J., et al. (2025). Artificial intelligence in undergraduate medical education: an updated scoping review. *BMC Medical Education*, 25(1), 1609.
11. Su, J., & Yang, W. (2023). Unlocking the Power of ChatGPT: A Framework for Applying Generative AI in Education. *ECNU Review of Education*, 6(3), 355-366. doi:10.1177/20965311231168423
12. Su, T., et al. (2025). The effectiveness of problem-based learning (PBL) in enhancing critical thinking skills in medical education: a systematic review and meta-analysis. *Frontiers in Education*, Volume 10 - 2025. doi:10.3389/feduc.2025.1565556
13. Tong, X., et al. (2025). The application of problem-based learning (PBL) guided by ChatGPT in clinical education in the Department of Nephrology. *BMC Medical Education*, 25(1), 1048. doi:10.1186/s12909-025-07427-w