ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

Learning with AI to prepare for working alongside AI

Hassan Khosravi
The University of Queensland

Generative AI is now woven into students' everyday learning, supporting tasks such as writing, coding, reasoning and analysis. Yet its educational value remains uncertain, because systems built to improve immediate task performance may weaken the cognitive and metacognitive processes on which durable learning depends. At the same time, there are strong calls for education to better prepare students for an AI-rich future of work and the capabilities it demands. This paper tackles this double challenge through the design of AI tools that prioritise learning over performance while developing the capabilities that matter for working with AI: creating original work, exercising sound judgement and achieving disciplinary mastery. We first outline the design framework we adapt for building such tools, in which AI learning companions are embedded within authentic learning environments. We then present RiPPLE, a co-creation platform used by more than 100,000 students, where AI companions support learners as they create learning resources, evaluate the work of peers, and engage in personalised practice, with each activity targeting one of the three capabilities. We conclude with reflections and selected empirical findings from the development and deployment of RiPPLE, drawing out lessons for designing AI tools that put learning first.

Implications for practice or policy:

- Institutions selecting AI tools should ask whether a tool is designed for learning or merely for task completion and evaluate it by what learners can do once AI support is withdrawn.
- Developers of educational AI should embed companions within learning environments, with persistent learner models, educator oversight and agency-preserving support, rather than deploying standalone chatbots.
- Educators can use co-creation activities supported by AI companions to develop innovation, judgement and disciplinary mastery alongside course content.

Keywords: artificial intelligence in education, generative AI, learning companions, co-creation, learner agency, future of work

Introduction

Generative AI (GenAI) has become part of the ordinary machinery of studying, called on to summarise readings, test reasoning, draft and debug, and analyse data (Chung et al., 2026). Institutional responses are maturing from bans towards structured adoption, yet students themselves are often unsure where productive use ends and self-defeating use begins, with some describing the worry of "cheating themselves" of understanding (Chung et al., 2026). The measurable news looks encouraging at first glance: meta-analyses find that students using ChatGPT and similar systems earn better marks on assignments and examinations (Deng et al., 2025). But those gains are recorded while the scaffolding is still in place, and the evidence base has been criticised as an effect still in search of a cause (Weidlich et al., 2025).

Beneath this uncertainty lies a distinction the learning sciences have long insisted on: performance during instruction is an unreliable guide to durable learning, and the conditions that inflate one frequently depress the other (Soderstrom & Bjork, 2015). GenAI turns this old distinction into a live paradox (Yan et al., 2025). In high school mathematics, students given an unrestricted AI assistant solved more practice problems while they had it, then scored significantly below unassisted peers once it was removed (Bastani et al., 2025). Essay quality rises even as learners hand the work of monitoring their own understanding to the tool, a pattern labelled metacognitive laziness (Fan et al., 2025), and sustained assistance can breed dependence, with self-regulated engagement dropping once support is withdrawn (Darvishi et al., 2024).

At the same time, a second pressure is mounting. AI is reshaping knowledge work itself, and education systems are being called on to prepare students for workplaces in which working with AI is routine (OECD, 2026). The uncomfortable observation is that the capabilities such calls emphasise, including the evaluative judgement and metacognitive regulation needed to appraise and direct fluent machine output (TEQSA, 2026), are precisely the capabilities that careless AI use in education erodes. Used as an answer engine, AI undermines the very capacities working with AI demands.

This paper treats these two challenges as one design problem: building AI tools that prioritise learning over performance while cultivating the capabilities that matter for working with AI. We first examine why general-purpose AI tools fall short of this goal and what working with AI actually demands of graduates. We then outline the design framework we adapt, in which AI learning companions are embedded within authentic learning environments, and present an in-depth illustrative design case of RiPPLE, a co-creation platform in which three AI companions support students as they create, evaluate and practise. We conclude with reflections, selected empirical findings, and lessons for designers, educators and institutions.

Background: learning with AI, working with AI

The learning–performance paradox

The GenAI now entering classrooms was not designed for them. Its design brief was productivity, meaning useful output, less effort and faster completion, and by that brief it succeeds: GenAI has been found to reduce writing time by 40% while raising output quality (Noy & Zhang, 2023), and to lift the productivity of customer support workers, with the largest gains for novices (Brynjolfsson et al., 2025). Education, however, inverts each element of that brief. Learning requires the learner to expend effort, errors are diagnostic opportunities rather than defects, and success is measured not by the artefact produced but by what the learner can later understand, transfer and do unaided (Soderstrom & Bjork, 2015). Table 1, below, summarises the contrast.

Table 1.
Comparison of AI designed for work and AI designed for learning

Dimension	AI for work (productivity)	AI for learning (capability)
Purpose	Complete tasks faster and to a higher standard	Develop the learner's own capability over time
Cognitive effort	Reduced, to maximise efficiency	Preserved and directed toward learning
Errors	Inefficiencies to be eliminated	Diagnostic signals and learning opportunities
Measure of success	Quality and efficiency of outputs	Retention, transfer and metacognitive growth
Key failure mode	Productivity gains that mask skill atrophy	Performance gains that mask absent learning

Retrofitting at the prompt level does not resolve the mismatch. Study modes and Socratic guardrails soften the harms of unstructured use (Bastani et al., 2025), but their measured effects on learning are often near null and unlikely to match the gains of purpose-built tutoring systems (Borchers & Shou, 2025) because the limits are architectural: prompted tutors are reactive rather than steering an instructional trajectory, stateless with respect to the learner's knowledge and goals, and bound to single tasks rather than a curriculum. Decades of research on educational agents point the same way. Learning companions, pedagogical agents, dialogue tutors and intelligent tutoring systems showed the value of peer-like roles, learner modelling, dialogic scaffolding and adaptive support (Biswas et al., 2005; Chan & Baskin, 1988; Chou et al., 2003; Graesser et al., 2005; VanLehn, 2011), but also that these strengths come from being designed into a learning environment, not bolted onto a general-purpose tool (Roll & Wylie, 2016).

Preparing students for a GenAI-integrated future

The pressure to respond is no longer only pedagogical but regulatory. In Australia, the Tertiary Education Quality and Standards Agency asks providers to demonstrate two things at once in a GenAI-integrated future: that learning has genuinely occurred despite ubiquitous AI assistance, and that students are developing the adaptive capabilities graduates will need in environments where human-AI collaboration is the norm (TEQSA, 2026). The resource defines four such capabilities: digital literacy, using digital and AI tools effectively, ethically and safely; distributed cognition, thinking productively across people, tools and technologies; hybrid metacognition, regulating one's thinking and learning within human-AI networks; and life-long learning. All four are underpinned by learner agency and regulation, and built on deep disciplinary knowledge as the epistemic foundation for authentic judgement (TEQSA, 2026). Similar expectations are emerging internationally (OECD, 2026).

These capabilities are precisely what the learning-performance paradox threatens. Hybrid metacognition, the planning, monitoring and adjusting of learning within human-AI systems, is exactly what learners offload when the tool does the thinking (Fan et al., 2025); deep disciplinary knowledge is bypassed when answers are delegated rather than constructed; and digital literacy in TEQSA's sense, deciding how, when and why to engage with GenAI rather than merely operating it, cannot develop through unreflective answer-seeking. Notably, the same resource endorses desirable difficulties and productive struggle as evidence-informed foundations, and calls for learning environments that make the processes of learning visible. The regulatory demand and the learning sciences therefore converge on one design requirement: tools in which students generate, judge and consolidate, rather than merely receive, serve assurance of learning and capability development at once.

An adapted design framework

Our design work adopts the account of AI learning companions proposed by Khosravi et al. (2026): an adaptive, pedagogically grounded and responsibly designed AI-mediated learning partner, embedded in an authentic learning environment, that supports durable learning across diverse learner populations. The definition carries eight design commitments, summarised in Figure 1 below: a companion is pedagogically grounded, adaptive, responsibly designed, present as a recognisable character, agency-preserving, embedded, learning-centred and inclusive. In practical terms, this means a companion needs a persistent model of the learner, a place inside the curriculum rather than beside it, oversight by educators, and a role that supports the learner's thinking without taking it over.

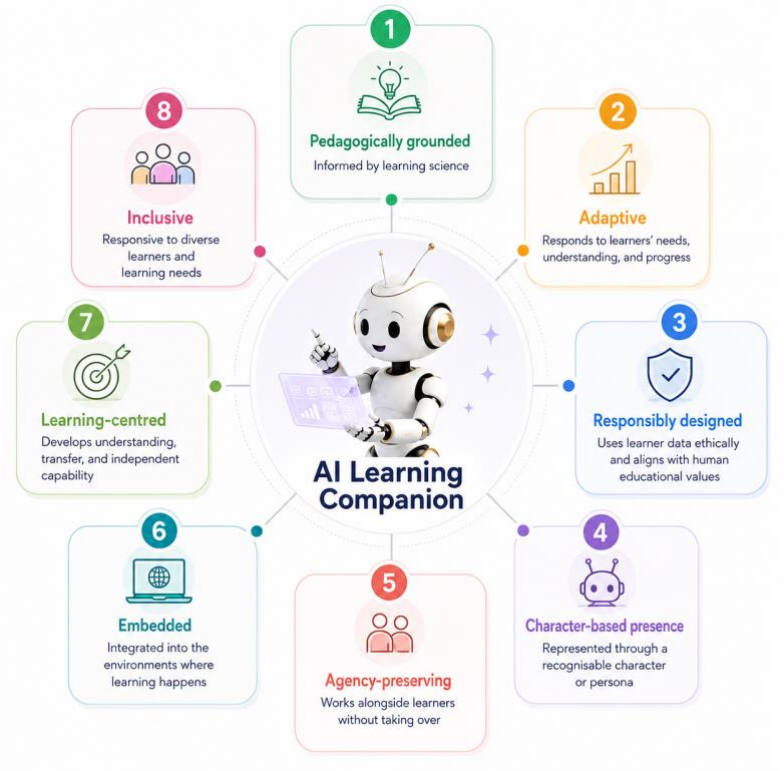


Figure 1. The AI learning companion and the eight design commitments it carries (adapted from Khosravi et al., 2026).

Our adaptation embeds these companions within a co-creation learning environment. Co-creation, in the students-as-partners tradition, positions students as contributors to the course's learning materials and processes rather than recipients of them (Bovill, 2020; Dollinger & Lodge, 2020), and recent work extends the partnership to AI, with GenAI supporting students during the creation of learning content rather than producing it for them (Pozdniakov et al., 2025). This pairing does deliberate work for the double challenge: creating resources exercises innovation, evaluating peers' work exercises judgement, and personalised practice consolidates mastery: precisely the generate–judge–consolidate pattern the previous section called for. The companions' role within each activity is constrained by the commitments above, particularly the preservation of learner agency.

An illustrative design case: RiPPLE

The co-creation ecosystem

RiPPLE is an adaptive co-creation platform, developed at a large public research university and now used in both higher education and schools (Khosravi, et al., 2019). Since its introduction it has been used by more than 100,000 students across a wide range of disciplines and course offerings, who together have created, evaluated and practised with a large shared repository of student-generated learning resources. The platform integrates with institutional learning management systems through LTI, single sign-on and grade export, so it operates inside existing course structures rather than alongside them. Its stated aim is to develop disciplinary mastery, sound judgement and the capacity to innovate, positioning students as contributors to a shared knowledge base rather than as recipients of content. These three target capabilities complement TEQSA's (2026) four adaptive capabilities: the three name what students develop through each activity, while the four name the capacities they exercise in doing so.

Figure 2, below, shows the ecosystem. The learner sits at the centre of a triadic cycle of activity: creating bite-sized learning resources such as questions, explanations and worked examples; evaluating the resources created by peers; and practising with resources recommended to them. Around this cycle, educators guide, moderate and configure; peers collaborate and provide feedback; and three AI learning companions, one per

activity, offer always-available support: Crea in creation, Modi in moderation and evaluation, and Praxie in practice. The platform continuously infers each learner's ability on every course topic from these activities and recommends practice aligned with their evolving mastery (Abdi et al., 2023).

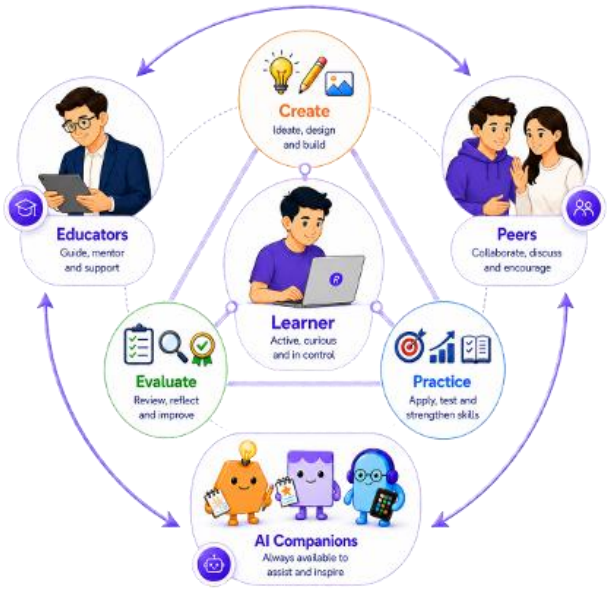


Figure 2. The RiPPLE co-creation ecosystem: the learner at the centre of a create–evaluate–practise cycle, supported by educators, peers and AI learning companions

Creating to innovate

Authoring a learning resource is demanding creative work: it requires explaining a concept, anticipating misconceptions and crafting plausible distractors; a form of elaboration associated with deeper processing (Hardy et al., 2014). Crea, the creation companion, supports this work in two modes, shown in Figure 3 below. While a student drafts a question, Crea provides dialogue support, checking the plausibility of distractors, the clarity of the stem and the alignment of the question with the intended level, while the student's own stem remains unchanged. Once the question is formally submitted, Crea returns structured feedback naming specific strengths and concrete improvements for the student to act on. The design stance throughout is that the companion critiques and suggests but never writes: the resource, and every decision about revising it, remains the student's. Deciding which of Crea's suggestions to accept, adapt or reject is itself an exercise in distributed cognition and digital literacy in TEQSA's (2026) sense. Analysis of more than 1,000 student-created questions and the corresponding AI feedback found the feedback broad in scope, well received by students for its clarity and tone, and, importantly, acted on: students used it to make concrete revisions to their work (Alsaiari et al., 2026).

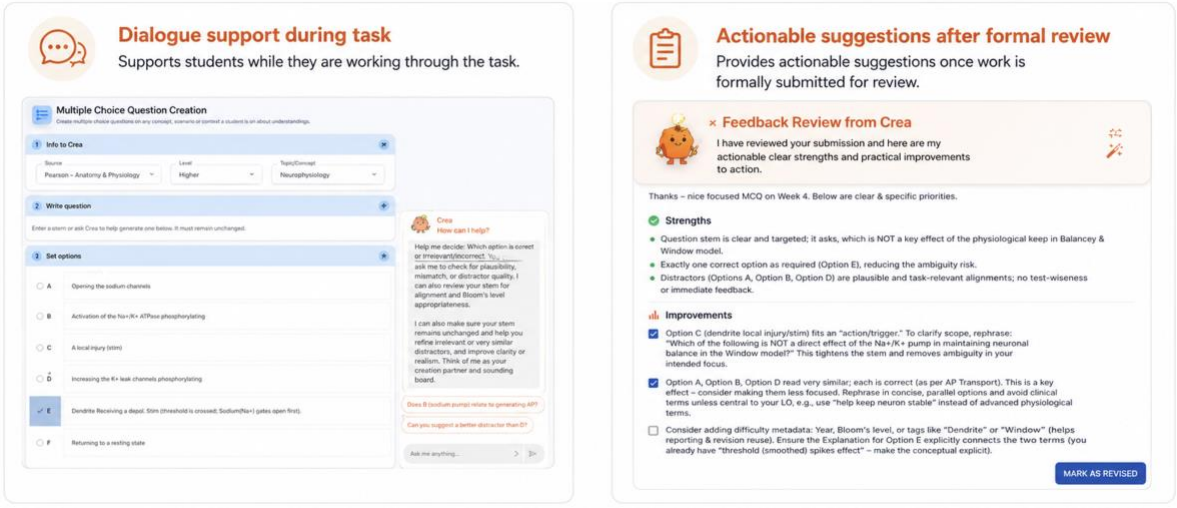


Figure 3. The creation companion in practice: dialogue support while a student creates a question (left) and actionable feedback once the work is submitted for review (right)

Evaluating to develop judgement

Every student-created resource is evaluated by multiple peers before it can reach the shared repository, and this evaluation is itself designed as an exercise in judgement: reviewers must assess correctness, clarity and pedagogical value, and articulate reasoning that the author can use (Topping, 2017). Modi, the moderation companion, scaffolds this work by analysing a student's draft feedback, surfacing missing rationale, and suggesting how comments could be made more specific and actionable, rather than writing the evaluation itself (Figure 4, below). The judgement, the rating and the final wording remain the reviewer's. In an environment where students will increasingly need to appraise fluent machine-generated content, the platform deliberately gives them structured, supported practice in appraising content of uncertain quality, with an AI companion coaching the quality of their judgement rather than substituting for it: structured development of the evaluative judgement that TEQSA (2026) identifies as central to appraising GenAI outputs.

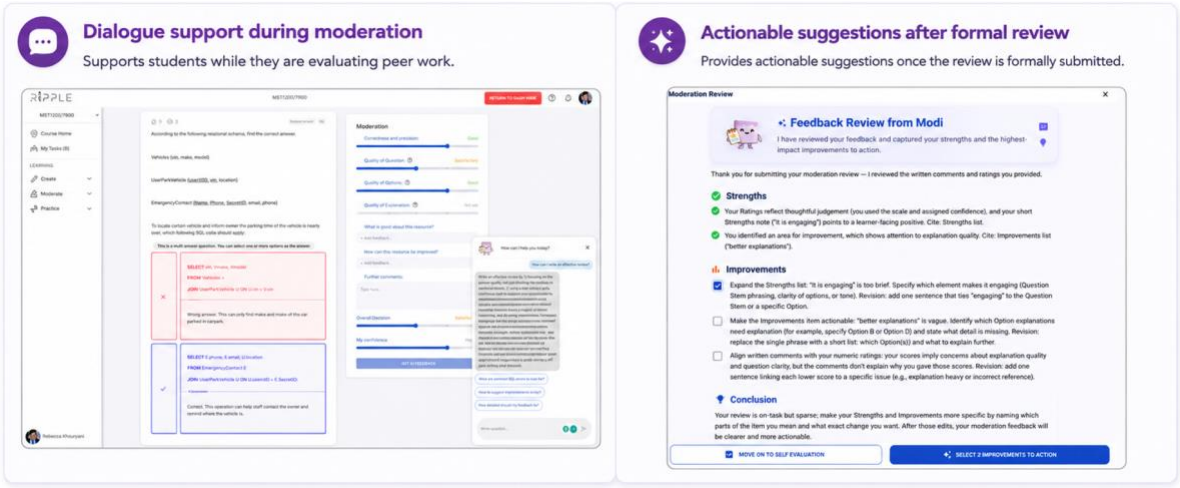


Figure 4. The moderation companion in practice: dialogue support while a student evaluates peer work (left) and review-oriented feedback once the moderation is submitted (right)

Practising for mastery

Practice in RiPPLE is organised as spaced retrieval, one of the most robust mechanisms for long-term retention (Roediger & Butler, 2011). A recommendation engine drives each learner's practice from persistent learner models built from performance, contribution history, peer evaluations and engagement traces (Khosravi et al., 2023). Central to this is an open learner model that estimates the student's ability on every course topic and makes those estimates visible to the learner (Figure 5, below, left), an approach associated with metacognitive

awareness and shared regulation (Bull, 2020). Recommendations are explainable in context, for example that an item was suggested because recent attempts show a recurring misconception (Khosravi et al., 2022). Praxie, the practice companion, draws on the same model: when a student seeks help during practice, including through voice-enabled interaction, the support offered is conditioned on what the model indicates the student already knows, asking for recall before offering hints (Figure 5, right). Because the models persist across the semester, support late in the course can build on misconceptions surfaced weeks earlier. The activity thus builds the deep disciplinary knowledge TEQSA (2026) treats as the foundation of every adaptive capability, while the open learner model makes hybrid metacognition tangible: students monitor, question and plan their learning within a human-AI system.

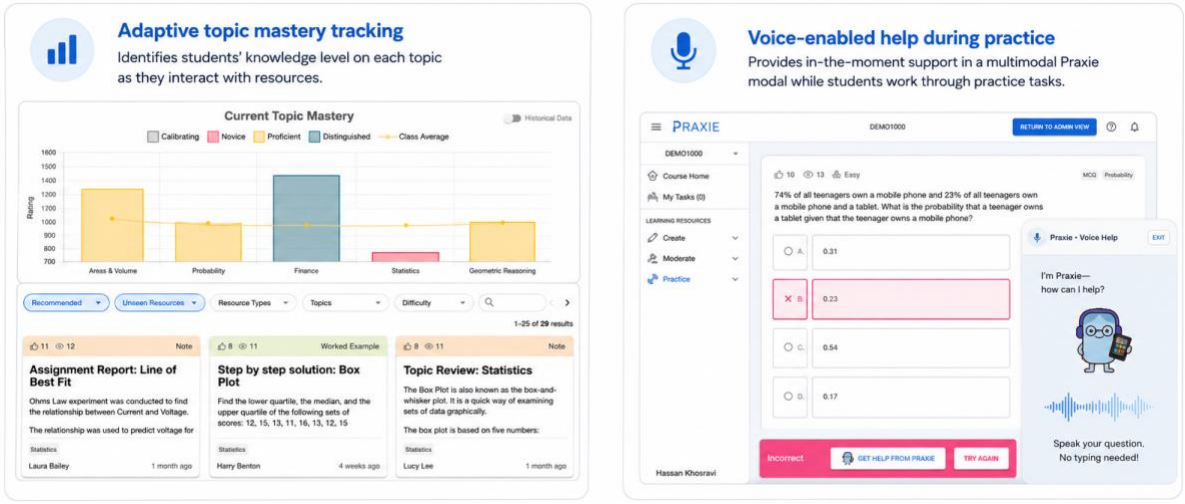


Figure 5. Practising for mastery in RiPPLE: the open learner model tracks each student's topic-level mastery (left) and informs the practice companion's in-the-moment, voice-enabled support (right)

The role of educators

None of the above operates without educators, and by design. Educators configure the activities and align them to weekly topics and assessment; they moderate what enters the shared repository, so publication remains human-gated even where companions have assisted; and they retain visibility into learner models, companion interactions and the health of the resource pool, with the ability to intervene where needed. The companions are therefore accountable to the course, not merely to the individual request: what counts as a good resource, a fair evaluation or appropriate support is ultimately set and supervised by the educator. This positions the educator not as a bystander to student–AI interaction but as the architect of the environment within which that interaction happens.

This oversight is itself AI-supported, in a way that concentrates rather than replaces expert judgement. When the peer review process is flagged by students or by the system as requiring expert human judgement; for example, where reviewer disagreement on a resource runs well above the course average, the case is queued for an educator. An AI review then independently examines the resource, the flags raised and the peer reviewers' feedback, and presents the educator with a structured summary: a judgement on the resource, a response to each flag, an appraisal of the quality of the peer reviews, and suggested feedback for improvement (Figure 6, below). The educator makes the final call, confirming, amending or overriding the AI's judgement, and decides the feedback that reaches the student. Expert human judgement stays at the decision point; the AI absorbs the preparatory work of each contested case. The arrangement is a working example of hybrid intelligent feedback, in which AI-generated and human feedback are responsibly combined to improve both the quality and the scalability of feedback (Banihashem et al., 2025).

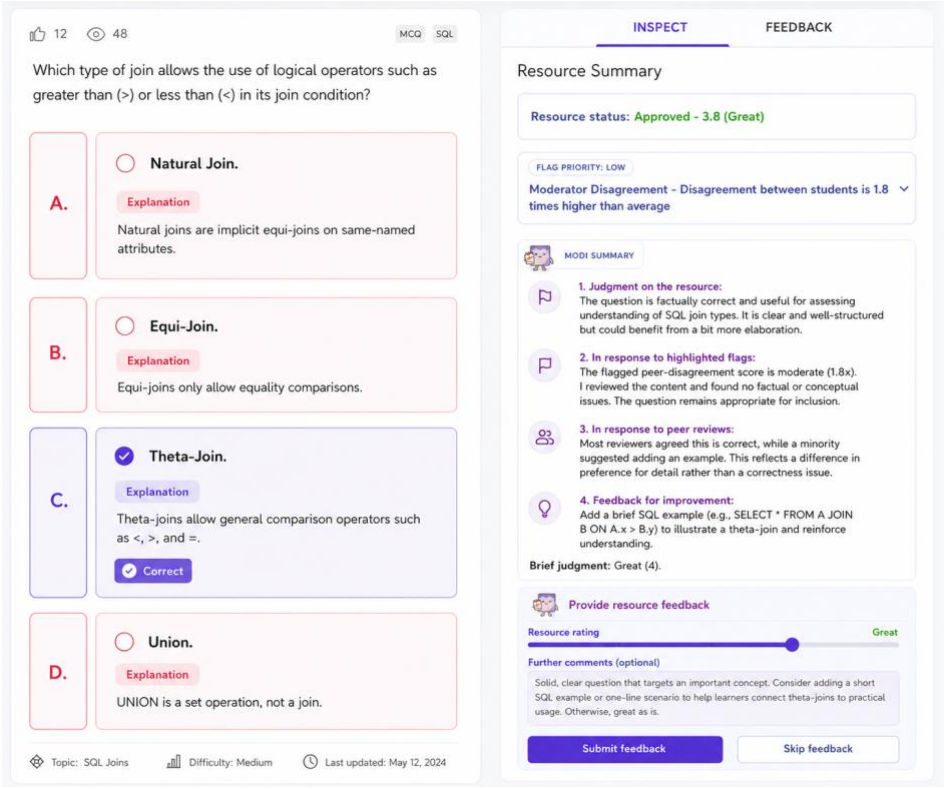


Figure 6. AI-supported educator moderation: a flagged resource (left) and the AI's independent review of the resource, flags and peer feedback, prepared for the educator's final judgement (right)

Table 2, below, summarises how each activity, companion and safeguard maps onto the three capabilities; the safeguards column is where the design commitments of Figure 1 become concrete.

Table 2.
Capabilities, activities, companions and safeguards in RiPPLE

Capability	Activity	Companion role	Key safeguards
Innovation	Creating learning resources	Crea: dialogue support while drafting; structured feedback on submission	Critiques, never writes; the resource remains the student's
Sound judgement	Evaluating peers' resources	Modi: analyses draft feedback, surfaces gaps, suggests improvements	Judgement, rating and wording remain the reviewer's
Disciplinary mastery	Personalised, spaced practice	Praxie: adaptive, voice-enabled help conditioned on the learner model	Open learner model; recall before hints; explainable recommendations

Discussion

For designers, the case yields four transferable moves, each answering a failure mode of general-purpose AI identified earlier. Where standalone LLMs are bound to the task in front of them, RiPPLE shows that the unit of design is the learning environment: its companions derive their value from activities, creating, evaluating and practising, that were chosen because they generate learning, and the same companions dropped into an environment with nothing to co-create would have little to offer but answers. Where general-purpose AI displaces the learner's effort, the platform positions companions inside work students are already doing, so they can question, critique and coach without taking the work over. Where prompted tutors are stateless, persistent and open learner models let support follow the learner's journey, and let learners see and steer it, rather than merely respond to requests. And where oversight by guideline depends on educator vigilance,

oversight built into the platform's mechanics, such as human-gated publication and AI-supported review of flagged cases, holds under the pressure of a busy semester. Together these are less a feature list than a design stance: AI belongs in the processes of learning, not at its outputs.

The empirical picture from the platform's development is encouraging but appropriately cautionary. Students engage with and value companion feedback during creation, and use it to make concrete revisions (Pozdniakov, 2025). At scale, that reception has held: as of 2026, deployment analytics show AI-supported feedback on resource creation rated 4.26 out of five on average across more than 70,000 instances, and feedback on peer reviews rated 4.24 across more than 300,000 instances. At the same time, studies on the platform found that sustained AI assistance can reduce students' independent engagement once support is withdrawn (Darvishi et al., 2024); a finding that directly shaped the agency-preserving design stance described above: companions suggest and critique, but do not produce, and support is designed to be reduced as competence grows. This is design evidence rather than outcome evidence: whether the three capabilities durably develop requires delayed, unaided measures of retention, transfer and judgement, which is the clear next step. Inclusive design across differences in language, ability and prior opportunity likewise remains ongoing work. One further institutional implication deserves note: because the platform records how students create, review and practise, it yields exactly the learning-process evidence that TEQSA (2026) asks institutions to build infrastructure for, so the same design that develops the capabilities also helps assure them.

The platform is also a site for refining what companion feedback should be. A semester-long randomised controlled trial in an adaptive educational platform compared directive, metacognitive and blended AI-generated feedback, and found the blend prompted the most revision while retaining the reflective benefits of metacognitive scaffolding (Alsaiari et al., 2026); a further randomised experiment found that emotionally enriched AI feedback reduced students' negative emotions towards feedback without compromising engagement or the quality of revised work (Alsaiari et al., 2025). Together these point to where companion support is heading: feedback calibrated not only in content but in form and tone, directive enough to enable action, metacognitive enough to build self-regulation, and emotionally attuned enough to sustain well-being.

The evaluation standard these limitations imply is the same one the paper began with. A tool that prioritises learning over performance must be judged by what learners can still do when the AI is withdrawn, and that is also the truest measure of their readiness to work with it. What students carry into an AI-rich workplace is not the output their tools produced, but the capability that remains: the ideas they can originate, the quality they can judge, and the knowledge they command.

Conclusion

Students are already using general-purpose AI, and it is neither realistic nor fair simply to ask them not to use it. The relevant choice is therefore not between AI and human support, but between generic systems designed primarily to provide answers and purpose-built tools designed to promote learning. This paper has argued that the learning–performance paradox and the demands of an AI-rich future of work are one design problem with one solution: learning environments with AI companions built in, where students create, evaluate and practise with support that preserves their agency and educators' oversight. RiPPLE offers one worked answer, and an invitation for more: the sooner learning with AI is designed deliberately, the better prepared students will be for a future of working with it.

References

Abdi, S., Khosravi, H., Sadiq, S., & Gašević, D. (2020). Complementing educational recommender systems with open learner models. In *Proceedings of the tenth international conference on learning analytics & knowledge* (pp. 360–365). ACM. https://doi.org/10.1145/3375462.3375520

Alsaiari, O., Baghaei, N., Lahza, H., Lodge, J. M., Boden, M., & Khosravi, H. (2025). Emotionally enriched AI-generated feedback: Supporting student well-being without compromising learning. *Computers & Education, 239*, 105363. https://doi.org/10.1016/j.compedu.2025.105363

Alsaiari, O., Baghaei, N., Lodge, J. M., Noroozi, O., Gašević, D., Boden, M., & Khosravi, H. (2026). Directive, metacognitive, or a blend of both? A comparison of AI-generated feedback types on student

engagement, confidence, and outcomes. *Computers and Education: Artificial Intelligence, 10*, Article 100553. https://doi.org/10.1016/j.caeai.2026.100553

Alsaiari, O., Baghaei, N., Lodge, J. M., Gašević, D., Winstone, N., & Khosravi, H. (2026). Making AI-generated feedback matter: A large-scale study of feedback workflows and student enactment. *arXiv preprint* arXiv:2608.11625.

Banihashem, S. K., Noroozi, O., Khosravi, H., Schunn, C. D., & Drachsler, H. (2025). Pedagogical framework for hybrid intelligent feedback. *Innovations in Education and Teaching International*. *63*(2), 554-570. https://doi.org/10.1080/14703297.2025.2499174

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences, 122*(26), e2422633122. https://doi.org/10.1073/pnas.2422633122

Biswas, G., Leelawong, K., Schwartz, D., & Vye, N. (2005). Learning by teaching: A new agent paradigm for educational software. *Applied Artificial Intelligence, 19*(3–4), 363–392. https://doi.org/10.1080/08839510590910200

Borchers, C., & Shou, T. (2025). Can large language models match tutoring system adaptivity? A benchmarking study. In *Artificial intelligence in education (AIED 2025)* (pp. 407–420). Springer. https://doi.org/10.1007/978-3-031-98417-4_29

Bovill, C. (2020). Co-creation in learning and teaching: The case for a whole-class approach in higher education. *Higher Education, 79*(6), 1023–1037. https://doi.org/10.1007/s10734-019-00453-w

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics, 140*(2), 889–942. https://doi.org/10.1093/qje/qjae044

Bull, S. (2020). There are open learner models about! *IEEE Transactions on Learning Technologies, 13*(2), 425–448. https://doi.org/10.1109/TLT.2020.2978473

Chan, T.-W., & Baskin, A. B. (1988). Studying with the prince: The computer as a learning companion. In *Proceedings of the International Conference on Intelligent Tutoring Systems* (pp. 194–200).

Chou, C.-Y., Chan, T.-W., & Lin, C.-J. (2003). Redefining the learning companion: The past, present, and future of educational agents. *Computers & Education, 40*(3), 255–269. https://doi.org/10.1016/S0360-1315(02)00130-6

Chung, J., Henderson, M., Slade, C., Liang, Y., Pepperell, N., Corbin, T., Walton, J., Yu, A. S., Bearman, M., Buckingham Shum, S., Fawns, T., McCluskey, T., McLean, J., Oberg, G., Seligmann, A., Shibani, A., Bakharia, A., Lim, L.-A., & Matthews, K. E. (2026). The use and usefulness of GenAI in higher education: Student experience and perspectives. *Computers and Education Open, 7*, 100347. https://doi.org/10.1016/j.caeo.2026.100347

Darvishi, A., Khosravi, H., Sadiq, S., Gašević, D., & Siemens, G. (2024). Impact of AI assistance on student agency. *Computers & Education, 210*, 104967. https://doi.org/10.1016/j.compedu.2023.104967

Deng, R., Jiang, M., Yu, X., Lu, Y., & Liu, S. (2025). Does ChatGPT enhance student learning? A systematic review and meta-analysis of experimental studies. *Computers & Education, 227*, 105224. https://doi.org/10.1016/j.compedu.2024.105224

Dollinger, M., & Lodge, J. (2020). Student-staff co-creation in higher education: An evidence-informed model to support future design and implementation. *Journal of Higher Education Policy and Management, 42*(5), 532–546. https://doi.org/10.1080/1360080X.2019.1663681

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

Graesser, A. C., Chipman, P., Haynes, B. C., & Olney, A. (2005). AutoTutor: An intelligent tutoring system with mixed-initiative dialogue. *IEEE Transactions on Education, 48*(4), 612–618. https://doi.org/10.1109/TE.2005.856149

Hardy, J., Bates, S. P., Casey, M. M., Galloway, K. W., Galloway, R. K., Kay, A. E., Kirsop, P., & McQueen, H. A. (2014). Student-generated content: Enhancing learning through sharing multiple-choice questions. *International Journal of Science Education, 36*(13), 2180–2194. https://doi.org/10.1080/09500693.2014.916831

Khosravi, H., Denny, P., Moore, S., & Stamper, J. (2023). Learnersourcing in the age of AI: Student, educator and machine partnerships for content creation. *Computers and Education: Artificial Intelligence, 5*, Article 100151. https://doi.org/10.1016/j.caeai.2023.100151

Khosravi, H., Gašević, D., Sadiq, S., Yan, L., Lodge, J., Tangen, J., Denny, P., DiCerbo, K., Buckingham Shum, S., & Baker, R. S. (2026). Building AI companions that prioritise learning over performance. *arXiv.* https://doi.org/10.48550/arXiv.2605.04816

Khosravi, H., Kitto, K., & Williams, J. J. (2019). RiPPLE: A crowdsourced adaptive platform for recommendation of learning activities. *Journal of Learning Analytics*, *6*(3), 91–105. https://doi.org/10.18608/jla.2019.63.12

Khosravi, H., Shum, S. B., Chen, G., Conati, C., Tsai, Y.-S., Kay, J., Knight, S., Martinez-Maldonado, R., Sadiq, S., & Gašević, D. (2022). Explainable artificial intelligence in education. *Computers and Education: Artificial Intelligence, 3,* Article 100074. https://doi.org/10.1016/j.caeai.2022.100074

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science, 381*(6654), 187–192. https://doi.org/10.1126/science.adh2586

OECD. (2026). *OECD digital education outlook 2026: Exploring effective uses of generative AI in education*. OECD Publishing. https://doi.org/10.1787/062a7394-en

Pozdniakov, S., Brazil, J., Mohammadi, M., Dollinger, M., Sadiq, S., & Khosravi, H. (2025). AI-assisted co-creation: Bridging skill gaps in student-generated content. *Journal of Learning Analytics*, *12*(1), 129-151. https://doi.org/10.18608/jla.2025.8601

Roediger, H. L., & Butler, A. C. (2011). The critical role of retrieval practice in long-term retention. *Trends in Cognitive Sciences, 15*(1), 20–27. https://doi.org/10.1016/j.tics.2010.09.003

Roll, I., & Wylie, R. (2016). Evolution and revolution in artificial intelligence in education. *International Journal of Artificial Intelligence in Education, 26*(2), 582–599. https://doi.org/10.1007/s40593-016-0110-3

Soderstrom, N. C., & Bjork, R. A. (2015). Learning versus performance: An integrative review. *Perspectives on Psychological Science, 10*(2), 176–199. https://doi.org/10.1177/1745691615569000

Tertiary Education Quality and Standards Agency. (2026). *Assuring quality learning in a gen AI-integrated future: The role of adaptive capabilities*. TEQSA. https://www.teqsa.gov.au/sites/default/files/2026-06/assuring-quality-learning-in-a-gen-AI-integrated-future.pdf

Topping, K. (2017). Peer assessment: Learning by judging and discussing the work of other learners. *Interdisciplinary Education and Psychology, 1*(1), 1–17. https://doi.org/10.31532/InterdiscipEducPsychol.1.1.007

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist, 46*(4), 197–221. https://doi.org/10.1080/00461520.2011.611369

Weidlich, J., Gašević, D., Drachsler, H., & Kirschner, P. A. (2025). ChatGPT in education: An effect in search of a cause. *Journal of Computer Assisted Learning, 41*, e70105. https://doi.org/10.1111/jcal.70105

Yan, L., Greiff, S., Lodge, J. M., & Gašević, D. (2025). Distinguishing performance gains from learning when using generative AI. *Nature Reviews Psychology, 4*, 435–436. https://doi.org/10.1038/s44159-025-00467-5

Khosravi, H. (2026). Learning with AI to prepare for working alongside AI. 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. 261–271). https://doi.org/10.65106/apubs.2026.3247

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.

© Khosravi, H. 2026