ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

Friction by design – desirable difficulty in AI co-learning

Daniel Laurence, Nikolai Alksnis

Victoria University

AI is only the latest in a succession of educational technologies that have made learning easier and more frictionless, which, counterintuitively, has coincided with a population-level decline in cognitive ability and poor learning outcomes. Studies in cognitive offloading and productive struggle suggest it may be affordances of information technology exacerbated by AI that contribute to these poor outcomes. Describing the relevance of cognitive offloading, cognitive load theory and the expertise reversal effect, this paper explores the cognitive mechanisms associated with technological assistance that have been shown to affect learning, and considers pathways to redesign learning in a way that focuses on useful forms of learning friction, understood as the effortful work that consolidates learning. The paper locates these cognitive mechanisms near the learner experience where self-regulation holds the most potential to effect change when combined with educator support. The question for educators, designers and students will continue to evolve: does experiencing friction enhance learning? A heuristic is offered as a conceptual model of a process for teachers and designers to support learners at the point where cognitive offloading occurs.

Implications for practice or policy:

- Research exploring the relationship between cognitive decline, poor learning outcomes and mechanisms that create forms of learning friction is outlined.
- Levels for interventions in policy, practice and design are explored with a focus on a self-regulated learning heuristic.

Keywords: cognitive offloading, desirable difficulties, productive struggle, generative AI, friction

Learning in a frictionless world

According to the best measures available in developed countries, “Multiple indicators now show stagnation or decline in literacy, numeracy, problem solving, creativity, and general cognitive performance among adolescents” (Horvath, 2026, p. 1). Successive waves of information technology have freed people from effort in the work of learning and knowing. The internet connected the world’s information, smart devices placed that information permanently at hand, and search engines relieved memory of the burden of retention (Sparrow et al., 2011). Social media and news feeds inundate learners with relevant information that predictively meets their interest profiles, and generative AI multiplies this progression in emergent ways. AI assists not only with the storage and retrieval of knowledge but with the reasoning, structuring and producing what were, until recently, necessarily human cognitive work. AI assistance no longer arrives only through intentional chat interfaces, it is becoming ambient in the ordinary tools of every discipline, including image editors, code completion inside development environments, and drafting assistance embedded in word processors. AI tutors were initially supposed to revolutionise learning in a positive way (Khan, 2023), but these claims quietened after Khan acknowledged studies showed no educational gains for most students using their AI tutors (Eames et al., 2026). Rigorous RCT studies have demonstrated that overreliance on AI can have a substantial negative effect on learning in a range of settings, learning activities and countries (OECD, 2026). Affirming Horvath’s (2026) indicators and suggested causes, large scale empirical studies from China also link cognitive offloading directly with decreased learning (Stromberg et al., 2026). With students graduating into workplaces where AI tools are standard, it is critical to help establish enduring behaviours in learners that enable them to stay cognitively engaged in the meaningful aspects of what they are doing. The task envisaged by Ferdosipour (2026) is therefore not so much to resist the removal of friction but to understand the mechanism, and then intentionally redesign where the friction in learning should be. This begins with articulating the specific mechanisms associated with the emerging shift in student learning.

The mechanisms of friction

While there is building evidence that cognitive offloading is a core mechanism responsible for a decline in educational outcomes, the point at which, and the type of task offloaded also appears critical (Lodge & Loble, 2026). The mechanism connecting offloaded memory to reduced learning is well established in cognitive science and can be described simply as: people retain what they are made to think hard about, not what they are merely shown (Willingham, 2009). Work on desirable difficulties and the testing effect (Roediger & Karpicke, 2006) shows that conditions which slow or impede performance during learning, effortful retrieval in particular, improve long-term retention relative to conditions that ease performance in the moment. Kapur's (2008) research on productive failure adds a sequencing principle: learners who attempt unsolved problems before instruction retain and transfer the resulting concept better than learners given the concept first, because the failed attempt activates prior knowledge that instruction subsequently attaches to. This has been demonstrated in research examining the use of AI in sequencing studies of cognitive offloading (Lodge & Loble, 2026).

This sequencing of learning is also the founding principle of long established Problem-Based Learning (Schmidt, 1983), in which unresolved problems are deliberately placed before instruction to activate prior knowledge and drive elaboration through group reasoning (Hmelo-Silver, 2004). However, conflicting arguments from Kirschner et al. (2006) additionally suggest that minimally guided instructional approaches, including Problem-Based Learning, can underperform for novices because insufficient prior knowledge can lead to cognitive overload rather than productive struggle. This nuance is represented perhaps in Vygotsky's classic 'Zone of Proximal Development' where learning is optimal when a task is sufficiently challenging, but not too hard for a learner. Any attempt to introduce friction to a learning task would need to ensure it lands in this zone and not be perceived as too hard.

The type of cognitive offloading is also a critical distinction. Offloading cognitive work to a tool may be beneficial when it removes extraneous load and frees limited working memory (Sweller et al., 2011) but becomes detrimental when it removes the work of thinking itself (Risko & Gilbert, 2016). While it may not hurt to use a calculator to solve a long-form statistical equation, if there was no thinking about what the calculation meant, it may harm learning. A simple, but powerful question is to ask: what would this task have taught me if I had done it myself? Complicating such a test is that the decision to offload is dependent on how much a learner already knows and their own speculative meta-cognitive evaluation of their personal competence. As such, learners would benefit from encouragement and support to think objectively about any decision to offload parts of a learning task.

Studies have demonstrated that unstructured reliance on generative AI can improve immediate task performance while weakening longer term learning (OECD, 2026; Yan et al., 2024). This has been termed the performance paradox, in which the work improves as the learning declines and the quality of the output conceals the absence of the capability. Learners are not well placed to detect the paradox in themselves, because rhetorically smooth AI output produces an illusion of competence that is incorrectly experienced as mastery (Lodge & Loble, 2026). Again, a learner's metacognitive evaluation of their level of mastery is revealed as the mechanism that could benefit from being bolstered in some way.

At the same time learners are being provided more helpful AI supports, they are also being saturated in more information than ever before. Cognitive load theory (Sweller et al., 2011) has guided learning design for some time, and while its primarily focus is reducing cognitive load, it also suggests that forms of learning friction can shift over time, in a finding known as the expertise reversal effect. Instructional supports and task structures that benefit novice learners lose their value, and can actively impair learning, as expertise grows (Kalyuga et al., 2003). Worked examples that relieve a novice's overloaded working memory become redundant processing for a learner who has consolidated the underlying schemas, so the now redundant guidance begins to obstruct learning and becomes a form of unhelpful friction. The effect is usually discussed as a problem for instructional sequencing, but its implication for the offloading is that if the value of support reverses as capability consolidates, then the boundary between beneficial and detrimental friction must adapt with the growing expertise of learners. Any fixed rules about the positive types of cognitive offloading, whether permissive or restrictive, will be wrong for some learners some of the time.

When these cognitive mechanisms that contribute to both helpful and unhelpful learning friction are explored together, the transitional, highly individual and personalised nature of such friction becomes apparent. While some standardised approaches may be applicable to entire cohorts of students with a similar level of starting skills and knowledge, the shifting nature of expertise can quickly make such broadly applied learning designs irrelevant and possibly even detrimental. So, while existing broad level approaches to reducing learning friction are initially helpful, a more sustained approach must also rely upon learners discerning for themselves where effortful friction should be embraced. Given also the false confidence in ability that AI can impart, the use of interruptions, such as self-diagnostic micro-reflection tools, would likely help to prompt learners to think authentically about how they are engaging with AI.

Levels of friction

Most institutions have established high-level systems to guide the mapping of broad competencies and learning outcomes to subject units, which has undoubtedly reduced unhelpful forms of learning friction. It is the next level down in individual subject units where the shifting redundancies and rough edges of a standardised design more often create the experience of friction for learners. Just as the needs of the workplace change over time, so too will aspects of a subject, and so the learning design process is a continuous re-evaluation of points of friction. Below the course and subject level sits the more highly variable student experience, where individual goals, capabilities, needs and preferences may not always align with the design that is offered. The course, subject and student levels all potentially contribute unhelpful friction to the learning experience and require differing types of interventions.

Table 1.
Levels at which friction works

Level	At stake	Friction interventions	Variability
Whole course	Accreditation, industry feedback and reputation	Policies, guides and curriculum mapping	Low
Individual subject unit	Relevance, coherence	Learning design, alignment	Medium
Student experience	Engagement and learning	Individual instruction and advice	High

While existing systems of curriculum and learning design work to reduce unhelpful friction at the course and subject unit levels, the student experience level can be far more variable, personal and appears to be the point where mechanisms like cognitive offloading and the cumulative effects of frictionless learning affect individual learners. At this student level, any intervention is generally required to be tailored directly to an individual's adapting needs, and is where individual instruction has been the primary intervention. Learning will always, ultimately, be an individual choice, the power that institutional interventions can have at this level is mostly limited to advice and creating a context for learning. The various providers of AI assistance have tried to introduce optional controls to their AI agents, like Microsoft Copilot's "Study and Learn Mode" which will quickly reveal if asked that it "is designed to make you think instead of just telling you answers." While a promising development, and one of the few touchpoints where mitigating the student experience of cognitive offloading might be possible, some have pointed out that it still acquiesces to provide answers very quickly (Ofgang, 2026) and the experience of being led to an answer already known revisits the expertise reversal effect. The temptation for students to reach for easily generated answers, writing and artefacts is too much for many, in a system that still rewards product over process. Again, a student's metacognitive awareness of learning is a key factor that determines if they cognitively offload appropriately.

Another approach is to provide overt metacognitive prompts for learners to self-regulate their own learning. A range of studies have demonstrated that prompting students to think before asking AI for answers can be an effective method of encouraging self-regulated learning (Darvishi et al., 2024). While this approach often requires curation to specific disciplines and capability levels, it is essentially the long standing Socratic teaching practice in digital form. Nelson-Le Gall (1981) distinguished instrumental help-seeking, seeking just enough support to solve a problem independently, from executive help-seeking, seeking to make the task disappear. Aleven et al. (2003) found the specificity of a help request signals which is occurring: a vague request ("I don't get this, just give me the answer") marks executive help-seeking; a request anchored to a specific point of difficulty marks instrumental help-seeking. To interrupt this kind of executive help seeking, it is proposed that

a simple heuristic may be adapted and used to remind students of the early friction that is required for learning and provides particularly adult learners with a gentle nudge without anyone over regulating or invigilating them. Try → Ask → Explain (TAE):

- Try first: attempt the task far enough at least to name a specific point of difficulty.
- Ask after: frame any AI query around a challenging element rather than the whole task.
- Explain last: when genuinely stuck ask AI to explain, or to process routine tasks.

By sequencing the trying phase first, this TAE heuristic encourages learners to experience friction and difficulty that leads to learning. Such difficulty can become disengaging for learners, so there must follow a clear point to ask AI for assistance, but on a specific, difficult element of the task. Only when students are genuinely stuck should they then ask for full explanations from AI, or if the task is routine. Encouraging the student as the decision-making agent for when and how they cognitively offload underpins this approach, and relies on students developing effective metacognition, judgement and asking themselves honestly: can experiencing this friction be useful?

One particular impact of how learning friction is shifting deserves to be considered with additional care. Lodge and Loble (2026) warn that unstructured AI use will widen existing equity gaps. Students with strong domain knowledge and metacognitive skill are more likely to offload beneficially and accelerate, while students without those foundations, disproportionately those already experiencing disadvantage, are likely to offload detrimentally and bypass the learning they most need. The ethical dilemma that unfolds without some form of design intervention is becoming apparent. A curriculum that leaves friction placement to individual student judgement delegates the decision to the learners least equipped to make that judgement. Lodge and Loble (2026) suggest that purposeful teaching interventions, perhaps like TAE, may provide disadvantaged learners the support needed to improve meta-cognition and decision making about their learning.

Conclusion

This paper describes a historical decline in population cognitive capacities that has coincided with the rise of frictionless information technologies. The resolution offered here is not a case against artificial intelligence in learning, but an argument that friction has become dislocated from where it does its work. While cognitive offloading is the most recently demonstrated mechanism associated with this decline, cognitive load theory details how volume of information can become overwhelming, and the expertise reversal effect explains that learner supports can grow to become yet another obstruction that overloads attention. While some approaches to address these challenges can, and have, occurred at the course and subject levels through policy and design interventions, the mechanisms at work operate largely on a student experience level that can best be addressed by encouraging students to develop their own self-regulated learning.

Evolving frictionless technologies will not distribute costs and benefits evenly: they will reward learners who already have the metacognitive and domain foundation to use them well and quietly withdraw the struggle that would have built that foundation in everyone else. While institutions must delegate decisions about cognitive offloading and friction to individual judgement, we can still signal the importance of trying before asking, particularly to our most disadvantaged learners. This is a conceptual paper, not an evaluation. Its claims about where friction belongs are design hypotheses, and the heuristic proposed here is offered for trial. What should not wait for further validation is the more basic reframing: the question is no longer whether AI assistance is used, but where, for this learner, at this stage, the friction still needs to be.

This paper is a conceptual piece synthesising existing literature, involves no primary data collection or human participants, and therefore requires no ethics approval.

References

Aleven, V., Stahl, E., Schworm, S., Fischer, F., & Wallace, R. (2003). Help seeking and help design in interactive learning environments. *Review of Educational Research, 73*(3), 277-320. https://doi.org/10.3102/00346543073003277

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

Eames, T., Brunskill, E., Yamkovenko, B., Weatherholtz, K., & Oreopoulos, P. (2026). Computer-assisted learning in the real world: How Khan Academy influences student math learning. *Proceedings of the National Academy of Sciences, 123*(1), Article e2507708123. https://doi.org/10.1073/pnas.2507708123

Ferdosipour, A. (2026, June 1). *Why AI needs Vygotsky: The case for AI-based intentional friction*. The Learning Guild. https://www.learningguild.com/articles/why-ai-needs-vygotsky-the-case-for-ai-based-intentional-friction

Hmelo-Silver, C. E. (2004). Problem-based learning: What and how do students learn? *Educational Psychology Review, 16*(3), 235–266. https://doi.org/10.1023/B:EDPR.0000034022.16470.f3

Horvath, J. C. (2026). *Written testimony of Dr. Jared Cooney Horvath before the U.S. Senate Committee on Commerce, Science, and Transportation. U.S. Senate Committee on Commerce, Science, and Transportation.* https://www.commerce.senate.gov/wp-content/uploads/media/doc/Horvath_Written Testimony.pdf

Kalyuga, S., Ayres, P., Chandler, P., & Sweller, J. (2003). The expertise reversal effect. *Educational Psychologist, 38*(1), 23–31. https://doi.org/10.1207/S15326985EP3801_4

Kapur, M. (2008). Productive failure. *Cognition and Instruction, 26*(3), 379–424. https://doi.org/10.1080/07370000802212669

Khan, S. (2023). *Brave new words: How AI will revolutionize education*. Viking.

Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. *Educational Psychologist, 41*(2), 75–86. https://doi.org/10.1207/s15326985ep4102_1

Lodge, J. M., & Loble, L. (2026). *Artificial intelligence, cognitive offloading and implications for education.* Australian Network for Quality Digital Education, University of Technology Sydney. https://doi.org/10.71741/4pyxmbnjaq.31302475

Nelson-Le Gall, S. (1981). Help-seeking: An understudied problem-solving skill in children. *Developmental Review, 1*(3), 224–246. https://doi.org/10.1016/0273-2297(81)90019-8

OECD. (2026). *Artificial intelligence and student learning outcomes: Evidence from randomised controlled trials*. OECD Publishing. https://doi.org/10.1787/062a7394-en

Ofgang, E. (2026, February 27). *Putting ChatGPT's study mode through its paces*. Edutopia. https://www.edutopia.org/article/putting-chatgpts-study-mode-through-its-paces

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences, 20*(9), 676–688. https://doi.org/10.1016/j.tics.2016.07.002

Roediger, H. L., III, & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science, 17*(3), 249–255. https://doi.org/10.1111/j.1467-9280.2006.01693.x

Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science, 333*(6043), 776–778. https://doi.org/10.1126/science.1207745

Stromberg, D., Lei, V., & Wu, Y. (2026, June 2). *The generative AI learning penalty: Evidence from Chinese secondary education*. SSRN. http://dx.doi.org/10.2139/ssrn.6868618

Sweller, J., Ayres, P., & Kalyuga, S. (2011). *Cognitive load theory*. Springer. https://doi.org/10.1007/978-1-4419-8126-4

Willingham, D. T. (2009). *Why don't students like school? A cognitive scientist answers questions about how the mind works and what it means for the classroom.* Jossey-Bass.

Yan, L., Greiff, S., Teuber, Z., & Luan, H. (2024). Promises and challenges of generative artificial intelligence for human learning. *Nature Human Behaviour, 8*(10), 1839–1850. https://doi.org/10.1038/s41562-024-02004-5

ASCILITE 2026

The Future is Connection:

Collaboration, Creativity and Compassion in a Digital Society

Laurence, D., & Alksnis, N. (2026). Friction by design – desirable difficulty in AI co-learning. 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. 689–694). https://doi.org/10.65106/apubs.2026.3312

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.

© Laurence, D., & Alksnis, N. 2026.