



THE SCIENCE OF

LEARNING AND GENERATIVE AI

A PRACTICAL GUIDE FOR SECONDARY TEACHERS

ABOUT THIS GUIDE

The science of learning and generative AI:
A guide for secondary teachers was
developed by researchers in the School
of Education at La Trobe University.

It brings established science of learning research to bear on a fast-moving question for schools: how generative AI can be used in ways that support, rather than shortcut, the thinking that learning requires. The guide is intended as a practical resource for secondary teachers, school leaders, and system leaders.

AUTHORS

Professor Miriam Tanti, Deputy Dean Commercial Growth, School of Education, La Trobe University

Dr Alexia Maddox, Senior Lecturer in Pedagogy and Education Futures, School of Education, La Trobe University

Dr Emma Phillips, Lecturer in Digital Technology Pedagogy, School of Education, La Trobe University

HOW TO CITE THIS GUIDE

Tanti, M., Maddox, A., & Phillips, E. (2026). *The Science of Learning and Generative AI: A guide for secondary teachers*. La Trobe University. [<https://doi.org/10.26181/32775462>]

COPYRIGHT AND LICENCE

© La Trobe University 2026.

With the exception of any logos, branding, and third-party material, this guide is licensed under a Creative Commons Attribution–ShareAlike 4.0 International (CC BY-NC-SA 4.0) licence. You are free to share and adapt the material – including for your own teaching context – provided you give appropriate credit and distribute any adaptations under the same licence. The full terms are available at creativecommons.org/licenses/by-nc-sa/4.0.

ACCESSIBILITY

This guide has been designed to meet WCAG 2.1 Level AA. Images and diagrams include text descriptions, body text meets minimum size requirements, and colour is not used as the only means of conveying information. If you encounter an accessibility barrier, please contact Miriam at m.tanti@latrobe.edu.au.

ACKNOWLEDGEMENTS

The authors thank Professor Penny Van Bergen, Professor Charlotte Forwood, and Dr Dara Tafazoli who assisted with the drafting. We are especially grateful to Distinguished Professor Pamela Snow, who generously gave her time and expertise in reviewing it. Her deep knowledge of the science of learning, and of how research evidence is best translated into classroom practice, sharpened the guide's arguments and strengthened its usefulness to teachers. Any errors that remain are the authors' own.

La Trobe University acknowledges the Traditional Owners of the lands on which this work was produced. We pay our respects to Elders past and present, and recognise their continuing connection to land, waters, and community.

ENQUIRIES

Enquiries about this guide can be directed to Professor Miriam Tanti at m.tanti@latrobe.edu.au.

INTRODUCTION

The science of learning brings together research from cognitive science, educational psychology, and neuroscience to explain how students actually learn – how they process information, build knowledge in complex and evolving schemas, develop understanding, and apply what they know over time.

These findings have moved well beyond theory. They now inform how effective teachers plan, instruct, and support learning in classrooms, for all students, including those who face systemic or contextual barriers to successful learning (e.g., socio-economic disadvantage, neurodivergence or disability, being an English-language learner, geographic remoteness, Aboriginal or Torres Strait Islander background).



A note on scope

This guide is written for teachers working with secondary students. This reflects both the developmental considerations that shape how students at different stages engage with new information and build knowledge, and the current age requirements of most generative AI (GenAI) tools available for educational use – many of which set a minimum age of 13, though terms of use vary by tool and some require parental consent for users under 18 (teachers should check the current terms for any tool they deploy). While upper primary students are increasingly likely to encounter GenAI outside school, the cognitive and motivational questions this raises are sufficiently complex that they sit outside the scope of this guide. Teachers working in primary contexts should exercise particular caution: the research on schema formation in younger learners, and emerging evidence on motivational dependency once students begin using AI tools, suggests that the case for GenAI use in primary settings requires its own careful treatment.

WHAT COGNITIVE PSYCHOLOGY RESEARCH TELLS US ABOUT HOW STUDENTS LEARN

One of the most important findings from decades of cognitive psychology research is that learning is shaped by how human memory works.

In the cognitive sciences, learning is generally defined as a lasting change in long-term memory: something is learned only when new knowledge has been organised and stored there well enough to be drawn on later (van Merriënboer & Sweller, 2005). While learning involves sensory and attentional processes that precede it, working memory is where conscious processing – and therefore the focus of instructional design – begins. Working memory is extremely limited – it can only hold and process (“work with”) a small amount of new information at once. When students encounter too much unfamiliar information at the same time, their working memory can quickly become overloaded, making it harder to understand and learn new information and skills (Cowan, 2001, 2010). Over time, through careful task analysis and design, practice, retrieval, and meaningful engagement, knowledge becomes organised into schemas held in long-term memory (Anderson et al., 1977; Roediger & Karpicke, 2006; Sweller, 2011). These schemas allow students to think more efficiently, manage complex ideas, and apply their understanding to new situations. Students can also add to and modify these schemas right across their lifespan, as they are exposed to additional, often more elaborate content.

“THE AIM OF INSTRUCTION SHOULD BE TO ACCUMULATE RAPIDLY SYSTEMATIZED, COHERENT KNOWLEDGE IN LONG-TERM MEMORY.”

(van Merriënboer & Sweller, 2005).

Research also shows that how deeply students engage with information matters as much as, if not more than, how often they encounter it (Craik & Lockhart, 1972). Surface engagement, such as passive re-reading or copying, produces weaker, less durable learning than deeper engagement like explaining, comparing, or applying, especially when these are combined with spaced retrieval practice. This is a finding with direct implications for how we think about generative AI (GenAI).

This understanding of how memory works has direct implications for teaching and learning. Because knowledge develops gradually – through attention, practice, retrieval, and feedback – effective teaching must manage how much students are asked to process at once, while gradually building more complex knowledge over time (Kirschner et al., 2006; Rosenshine, 2012; Sweller et al., 2019). Learning is therefore more than accessing information or completing tasks. Students need

opportunities to actively think about, organise, apply, and reflect on knowledge in ways that genuinely strengthen understanding.

Not all cognitive effort works the same way. Some forms of challenge – retrieving information from memory, spacing practice over time, or mixing up different types of problems – actually improve long-term learning, even when they feel harder in the moment (Bjork & Bjork, 2011). This kind of desirable difficulty builds durable knowledge. So, the goal of good teaching is not to make learning easier, but to remove the cognitive clutter (i.e., to manage cognitive load) that wastes effort, so that students’ thinking is directed toward the processes that actually build lasting understanding: retrieving, elaborating, connecting, and applying (Sweller et al., 2019).

Humans have always used tools to support thinking and learning, from spoken language and writing to diagrams, calculators, books, and digital technologies. The tools students use shape how they think, not just what they can access (Hutchins, 1995). GenAI, however, is a different class of technology from the tools that preceded it. Earlier technologies such as search engines, calculators, and word processors helped students find, store, or display information, but the thinking still had to be done by the learner. Reading, reasoning, drafting and explaining were processes students had to do themselves. GenAI – whether accessed through a conversational chatbot like ChatGPT or Gemini,

a purpose-designed curriculum agent, or AI features embedded in platforms students already use – can now do those things too. It can produce explanations, construct arguments, write responses, and adapt to what a student asks in real time. This means a student can now obtain the products of reasoning, drafting, and explaining without doing the reasoning, drafting, or explaining themselves. GenAI does not just provide access to information; it can supply the output that the thinking would have produced. The risk is not that GenAI does the cognitive work that builds knowledge (it cannot); it is that GenAI removes the need for the student to do it, and that work is where the learning happens.

This is what makes GenAI both educationally significant and educationally challenging in ways that earlier technologies were not.

Some uses of GenAI can potentially deepen learning by supporting students to explain their thinking, explore ideas, get feedback, and reflect on what they already know and/or are learning. Other uses can work against learning by encouraging students to simply accept generated responses without productive cognitive effort or critical engagement, leading to shallow engagement or the illusion of understanding without genuine knowledge acquisition/schema development.

This guide explores that challenge through a practical framework grounded in the science of learning. Throughout, classroom examples and strategies are connected to key evidence-informed principles including Cognitive Load Theory (CLT; Sweller, 1988; Sweller et al., 2019), schema development, retrieval practice, explicit instruction, transfer, and self-regulated learning. Rather than offering fixed rules about AI use, the guide supports teachers to think carefully about when and how GenAI amplifies students' thinking and when it risks bypassing the sometimes-effortful learning it is meant to support. A note on the evidence base: the cognitive science this guide draws on is well established, but its application to GenAI specifically is necessarily more provisional. The classroom claims that follow are grounded in that research and in reasoned extrapolation from it; the direct empirical evidence on GenAI's effects on learning is still emerging, and the guide should be read in that light.

**THE EDUCATIONAL
VALUE OF GENAI
THEREFORE LIES NOT
IN THE TECHNOLOGY
ITSELF, BUT IN
WHETHER IT SUPPORTS
THE COGNITIVE EFFORT
REQUIRED FOR REAL
LEARNING - OR QUIETLY
UNDERMINES IT.**

KEY AREAS FOR LEARNING WITH GENAI

This guide focuses on four areas where GenAI and learning intersect most directly:

01

**ATTENTION AND
COGNITIVE LOAD**



02

**MEMORY AND
KNOWLEDGE**



03

**TRANSFER AND
PROBLEM-SOLVING**



04

**METACOGNITION AND
SELF-REGULATION**



These areas are closely connected. What happens in one area shapes what is possible in the others. Together, they provide a practical framework for thinking about how GenAI can amplify students' thinking, and where it can work against learning. The framework on the next page brings this together at a glance: how learning works, the four focus areas and how they connect, and where GenAI supports learning or gets in the way.

THE SCIENCE OF LEARNING + GENAI

A PRACTICAL FRAMEWORK FOR TEACHERS

THE GENAI CHALLENGE

GenAI can perform the very thinking through which learning happens...


it can support that thinking — or do it for students. Its value depends entirely on how the task is designed and implemented

GenAI undermines learning when...

it performs the cognitive work — drafting, reasoning, explaining — in place of the student

GenAI supports learning when...

it requires students to explain, retrieve, compare and construct — not just receive



Key Area 4 (metacognition and self-regulation) is drawn with a dashed rather than solid border to show it runs through and supports the other three areas, rather than standing apart from them.

Sweller et al. (2019) · Bjork & Bjork (2011) · Roediger & Karpicke (2006) · Craik & Lockhart (1972)

WHAT IS GENERATIVE AI?

Generative AI (GenAI) refers to a class of AI technology that produces new content – text, images, code, or speech – in response to a user’s input. Unlike earlier digital tools that retrieved or displayed existing information, GenAI generates responses by drawing on statistical patterns learned from large amounts of training data. It can explain concepts, answer questions, write drafts, provide feedback, and adapt its responses in real time based on what a student asks.

In classrooms, GenAI is accessed through different kinds of interfaces. Most commonly, this is a conversational chatbot – a tool students interact with through typed prompts and responses, such as ChatGPT, Microsoft Copilot, or Google Gemini. It also appears as purpose-designed agents built around specific curriculum content, and as AI features embedded within existing platforms students already use, such as Copilot in Microsoft Word or Teams, or AI writing tools in Google Docs. Whatever form it takes, the defining characteristic is the same: GenAI does not just retrieve information. It generates language, constructs arguments, and produces explanations – performing the kinds of cognitive work that, when done by the student, builds lasting understanding.

WHAT YOU NEED TO KNOW

It is important to understand what GenAI does when it generates a response, because this shapes everything about how it should be used in learning contexts. Despite the name, these systems are not intelligent in any meaningful cognitive sense – they do not think, reason, understand, or know in the way humans do. This distinction matters particularly for a guide like this one. The science of learning draws on cognitive science – research into how human minds process information, build knowledge, and develop understanding. That is the kind of learning process this guide is concerned with: the real cognitive development of students. GenAI, despite the label of intelligence, operates on an entirely different basis. It does not develop, it does not learn in the way students learn, and it does not understand what it produces. Using cognitive science

to support students’ use of these tools is not a contradiction – it is precisely the point. The science of human learning is what allows us to judge when GenAI supports genuine cognitive development and when it bypasses it entirely.

HOW IT WORKS

What GenAI systems do is predict: given everything that has come before in a sequence, a large language model (LLM) calculates the statistically most likely next token. The output can be fluent, coherent, and convincing – because human language is deeply patterned – but fluency is not comprehension. The model has no human-like access to meaning; it operates over learned patterns in data. This is why GenAI can produce confident errors – which some describe as false positives rather than hallucinations: outputs that match the statistical pattern of a correct answer without any mechanism for verifying that



they are correct. It has no reliable mechanism for recognising when it is wrong. Techniques such as retrieval-augmented generation (RAG) can reduce the frequency of these errors by grounding the model's output in verified sources, but they cannot remove them entirely. The reason is architectural, not incidental: the model is still selecting the most statistically likely continuation, with no independent way to check that continuation against the truth. Better tools will make errors rarer – but they will not change the underlying mechanism.

WHAT THIS MEANS FOR CLASSROOM USE

Because GenAI cannot verify its own outputs, students must. Using GenAI productively requires students to approach its responses with critical judgement – checking claims, questioning sources, and not accepting generated content at face value. These are not simply good habits for AI use. They are core information literacy skills, and they need to be explicitly taught alongside any classroom use of the tool. These skills are threaded

through the classroom examples that follow, grounded in practice rather than abstraction. What makes all of this harder is the fact that students are better placed to evaluate AI output critically when they already have background knowledge about a topic – but in many learning contexts, students are encountering content for the first time, without the prior knowledge needed to recognise what is missing or wrong. GenAI, used well, can prompt exactly this kind of thinking. Used without it, it can quietly undermine it.

1. MANAGING ATTENTION AND COGNITIVE LOAD WITH GENAI

Attention is where learning starts. Before students can process, understand, or remember new information, they first need to notice and focus on it (Willingham, 2009).

This sounds simple, but in practice it is one of the most significant challenges students (and by extension, teachers) face – particularly in digital learning environments where students are navigating multiple sources of information, notifications, and competing demands all at once.

GenAI adds another layer to this challenge. Used well, it can potentially help direct students' attention toward what matters most. Used poorly, it can potentially contribute to distraction, information overload, or fragmented thinking that works against learning.

HOW WORKING MEMORY SHAPES WHAT STUDENTS CAN LEARN

Attention and working memory are closely linked. As established in the previous section, humans' working memory is limited – students can only work with a small amount of new information at a time, especially when the material is unfamiliar and/or complex (Cowan, 2001, 2010). Attention acts as a gatekeeper for working memory: information that students actively attend to is selected for further processing, while information that is not attended to is quickly lost. When too much information competes for a student's attention at once, their working memory can become overloaded, making it hard to process new ideas and slowing or disrupting learning. This process varies depending on students' developmental stage and prior knowledge – considerations that are particularly relevant for teachers working across different year levels. This guide focuses on secondary students, broadly aligned with the age ranges for which current GenAI tools are designed for educational use.

Cognitive Load Theory (CLT) explains how these attentional and working memory limitations influence learning and why the way teachers design and sequence instruction matters so much (Sweller et al., 2019). Like all cognitive science research, CLT findings are largely derived from controlled experimental settings, and translating them to real classrooms requires professional judgement. The principles are robust guides, not rigid rules – they tell us what to pay attention to, not exactly what to do.

CLT draws a useful distinction between two types of cognitive load:

- **Intrinsic load** – the inherent complexity of what is being learned. Some topics are harder than others at different year levels, and that complexity is real.
- **Extraneous load** – unnecessary cognitive effort created by poor design, distractions, cluttered explanations, or too much information arriving at once.

Effective teaching aims to manage both intrinsic and extraneous cognitive load – not to make learning effortless, but to ensure students' effort is directed toward understanding the material rather than managing avoidable confusion. As noted in the previous section, some effort is genuinely desirable: retrieving information from memory, working through a challenging problem, or explaining an idea in your own words all involve effort that strengthens learning (Bjork & Bjork, 2011). The goal is to protect and direct that kind of effort, not eliminate it.

This distinction matters directly when thinking about GenAI.

A tool that helps students manage extraneous cognitive load – by clarifying a confusing explanation, breaking a complex task into steps, or providing a relevant example – can genuinely support learning. A tool that removes intrinsic cognitive effort by doing the thinking, reasoning, or problem-solving for the student undermines it. This is the distinction Lodge and Loble (2026) frame as beneficial versus detrimental cognitive offloading: using AI to reduce unnecessary load frees working memory for the thinking that matters, whereas using it to bypass that thinking altogether erodes the learning it is meant to support.

WHAT THIS MEANS FOR GENAI IN THE CLASSROOM

GenAI has the potential to support attention and reduce extraneous cognitive load. It can potentially provide scaffolds, titrated worked examples, simplified explanations, visual representations, and timely feedback – supports that help students focus their effort on understanding key ideas rather than managing unnecessary complexity. For students who are newer to a topic, this kind of targeted support can make the difference between desirable difficulty and unproductive overload.

At the same time, GenAI can potentially work against learning if it is not used with clear purpose. Long, dense AI-generated responses can overwhelm rather than clarify. Moving rapidly between prompts, generated answers, and related content can potentially fragment attention and prevent the sustained thinking that builds understanding. And when GenAI does too much of the cognitive work (e.g., drafting, explaining, reasoning), students may appear to be engaging with ideas while actually bypassing the thinking that builds lasting knowledge.

The following classroom example is a constructed vignette – a realistic but deliberately designed scenario that illustrates the above principles in action. Real classrooms are messier – students resist, misunderstand, and surprise.

The example shows what deliberate, principle-informed teaching can look like, not a guarantee of what will happen, but a model of what to aim for. This framing is applied to all four classroom examples in this guide. The chatbot responses shown here, and in the vignettes that follow, also represent best-case interactions. Real tools – especially general-purpose chatbots – often respond less precisely: giving longer, less focused answers, or simply supplying information rather than prompting thinking. The exchanges model what good task design and prompting make more likely, not what a tool will reliably do unprompted.

CLASSROOM EXAMPLE: YEAR 10 SCIENCE

MANAGING ATTENTION AND COGNITIVE LOAD WITH AN AI CHATBOT

BEFORE THE LESSON

Ms Orlando's Year 10 Science class is partway through a unit on climate systems. Tomorrow's lesson will introduce students to the enhanced greenhouse effect using an AI chatbot as a research and thinking tool.

Before students arrive, Ms Orlando makes two deliberate decisions about how the lesson will begin:

1. She plans a short review of prior knowledge. Students have already covered the water cycle and basic atmospheric science. She wants to activate what they already know before introducing new content – not as a warm-up ritual, but because bringing existing knowledge to mind reduces the gap between what students already understand and what they are about to learn.
2. She prepares a starting prompt for the chatbot. She knows from experience that open-ended access without a clear entry point tends to produce one of two problems: students who ask for everything at once and get overwhelmed, or students who don't know where to start and disengage. The starting prompt is designed to give every student a manageable, focused way in, and guides their focus towards the learning intention when using the chatbot.

SCIENCE OF LEARNING ACTIVATING PRIOR KNOWLEDGE



When students bring existing knowledge to the surface before encountering new content, new information has somewhere to connect. Without that activation, new content arrives in working memory with no existing structure to attach to – making it harder to process, organise, and retain. Ms Orlando's two decisions – the review and the starting prompt – are both attempts to manage this before the lesson begins.

THE LESSON BEGINS

Ms Orlando opens the lesson without the chatbot. She puts three questions on the board:

Review Questions – Whole class discussion

1. What do you already know about how Earth's atmosphere works?
2. What is the greenhouse effect, in your own words?
3. What do you think 'enhanced' might mean in the phrase 'enhanced greenhouse effect'?

The discussion runs for five minutes. Students share what they know. Ms Orlando listens carefully, noting where understanding is confident and where gaps appear. She does not correct every misconception immediately – she notes them, knowing the lesson will address them. By the end of the discussion, the class has a shared, if basic, foundation to work from.

SCIENCE OF LEARNING SCHEMA ACTIVATION AND WORKING MEMORY



Working memory works more efficiently when it can draw on knowledge already held in long-term memory. The brief review does exactly this – it surfaces what students already know, making it available as a foundation for what comes next. Students who begin a complex topic with their prior knowledge activated are better placed to connect new information meaningfully rather than processing it in isolation. The discussion also gives Ms Orlando real-time information about where the class is – which she will use to pitch the chatbot task at the right level of challenge, and when she introduces the lesson's learning intention and success criteria, which are displayed throughout the lesson to focus attention.

MODELLING THE SKILL

After introducing the learning intention and success criteria for today's lesson and before students open the chatbot, Ms Orlando takes ten minutes to model how to use it for learning. She projects her screen and works through an exchange with the class watching.

She begins by giving the chatbot a short set of standing instructions. These do not concern any single question; they shape every reply that follows. It is a deliberate move to reduce extraneous load. Setting the rules up front means the chatbot's answers support attention rather than overwhelming students, even before the class has asked anything.

STANDING INSTRUCTIONS FOR CHATBOT

Answer at a Year 10 level.
Use simple, clear language, no jargon.
Keep the explanation short.
Include one example or analogy to help understanding.

MS ORLANDO

"I'm going to show you how I'd approach this – not just what to type, but how I'm thinking about it. Watch what I do when the chatbot gives me more than I need, and what I do when I want to go deeper on one idea."

She types a deliberately focused first prompt and reads the response aloud, narrating her thinking as she goes.

MODELLED EXCHANGE

What is the difference between the natural greenhouse effect and the enhanced greenhouse effect? Explain just that one distinction.

The natural greenhouse effect is the normal warming of Earth caused by naturally occurring greenhouse gases that trap some heat in the atmosphere. The enhanced greenhouse effect occurs when human activities increase the amount of greenhouse gases, causing more heat to be trapped than would occur naturally. The key distinction is that the natural greenhouse effect is a natural process, while the enhanced greenhouse effect is caused by additional greenhouse gases released by humans.

MS ORLANDO

"Before I read anything else, I'm going to say it back in my own words: the natural effect is the normal warming of the Earth that makes it liveable, and the enhanced greenhouse effect is caused by human emissions that make it worse.

What do I want to understand next? I'm not going to ask for everything at once. I'm going to pick one thing I'm curious or unsure about from that response."

MS ORLANDO'S SCREEN – FOLLOW-UP

Why does adding more greenhouse gases cause more heat to be trapped in the atmosphere?

Adding more greenhouse gases causes more heat to be trapped because greenhouse gases absorb and re-radiate heat energy that Earth releases into the atmosphere. When there are more of these gases, a greater amount of heat is absorbed and sent back towards Earth's surface instead of escaping into space. This increases the amount of heat retained in the atmosphere, causing temperatures to rise.

A simple analogy is adding extra layers of insulation to a house—the more insulation there is, the harder it is for heat to escape.

MS ORLANDO

"The insulation in a house analogy is useful. Let me write that down in my own words before I move on – not copy it but say it in a way that makes sense to me, using full sentences. That's the habit I want you to build today: read, think, then say it yourself before moving on."

SCIENCE OF LEARNING MANAGING COGNITIVE LOAD THROUGH MODELLING



Using an AI chatbot for genuine learning is a new skill – and new skills need to be taught explicitly; they are not just about access to tools. When Ms Orlando thinks aloud, she makes visible to novices the decisions that experienced learners (experts) make automatically: choosing a focused question, resisting the pull to ask for everything, pausing to process before moving on. Students who have not yet developed these habits cannot be expected to apply them independently without first seeing them in action. Modelling also keeps cognitive load manageable – students are watching one exchange at a time, rather than navigating their own chatbot while simultaneously trying to figure out how to use it.

GUIDED PRACTICE

Ms Orlando now gives the class a starting prompt and asks them to work through it before moving to their own questions. She projects it and reads it aloud.

STARTING PROMPT FOR CHATBOT – PROVIDED TO ALL STUDENTS

We have just looked at the difference between the natural and enhanced greenhouse effect. Now choose one cause of the enhanced greenhouse effect and explain how it contributes to the problem.

Explicit Instruction – Provided to all students

You are going to try this prompt with the chatbot yourself. Once you can explain the chatbot's answer in your own words, decide on one follow-up question of your own.

Students open their chatbots. After three minutes, Ms Orlando pauses the class.

MS ORLANDO

"Before you go any further, turn to your partner. Student A, explain in your own words what that first response was saying. Student B, add one idea they left out. Don't read from the screen. Just say it. After each member of the pair has said one thing, we will then share a few responses with the class."

Pairs talk briefly. Ms Orlando moves around the room, listening. She hears mostly solid understanding, and one persistent misconception: several students believe the enhanced greenhouse effect is a recent and sudden phenomenon rather than a gradual intensification of a natural process. She asks a few students to share their responses (think-pair-share). She addresses it directly with the whole class before students continue.

SCIENCE OF LEARNING RETRIEVAL AND MISCONCEPTION CHECKING



Pausing after the first exchange – before students branch into their own questions – is a deliberate attempt to check that the foundational understanding is in place before new knowledge is added to it. If students carry a misconception forward, every subsequent piece of information they encounter will connect to a flawed structure. The brief pair discussion is also a low-stakes retrieval moment: students are asked to recall and articulate, not just re-read. This small act of retrieval begins to consolidate understanding before the lesson moves on and sharing it with the class may improve engagement in the task.

INDEPENDENT PRACTICE

Students now work independently, each pursuing their own follow-up questions. Ms Orlando moves between tables, watching screens and asking questions. She is not monitoring for rule-following – she is watching for the cognitive patterns that signal whether the chatbot is supporting or replacing thinking.

Two contrasting approaches emerge.

JORDAN'S SCREEN

Jordan

We have just looked at the difference between the natural and enhanced greenhouse effect. Now choose one cause of the enhanced greenhouse effect and explain how it contributes to the problem.

AI chatbot

One major cause of the **enhanced greenhouse effect** is the burning of **fossil fuels** such as coal, oil, and natural gas.

When fossil fuels are burned to produce electricity, power vehicles, or run industries, they release large amounts of **carbon dioxide (CO₂)** into the atmosphere. Carbon dioxide is a greenhouse gas that traps heat from the Sun. As more CO₂ builds up in the atmosphere, more heat is trapped, causing the Earth's average temperature to rise.

This extra warming strengthens the natural greenhouse effect, leading to problems such as **climate change, rising sea levels, more extreme weather events, and changes to ecosystems**.

In summary, burning fossil fuels contributes to the enhanced greenhouse effect by increasing the amount of carbon dioxide in the atmosphere, which traps more heat and warms the planet.

Jordan

what are the causes of climate change and what are countries doing about it.

AI chatbot

Climate change is mainly caused by human activities that increase greenhouse gases in the atmosphere, which trap heat and warm the Earth. The main causes include burning fossil fuels such as coal, oil, and gas for energy and transport, which releases large amounts of carbon dioxide; deforestation, which reduces the number of trees that absorb carbon dioxide and releases stored carbon; and agriculture and industry, which produce gases like methane and nitrous oxide. Although natural factors such as volcanic eruptions and changes in solar energy can influence climate, they occur much more slowly and do not explain the rapid warming seen today. Countries are responding to climate change by working together through international agreements such as the Paris Agreement, which aims to limit global warming, and by setting targets to reduce emissions. They are also taking action within their own countries by investing in renewable energy like solar and wind power, improving energy efficiency, protecting forests, and promoting cleaner transport such as electric vehicles. In addition, governments are adapting to the effects of climate change by building flood defenses, improving disaster preparedness, and developing ways to cope with extreme weather.

Jordan is scrolling through the response, highlighting text. When Ms Orlando asks what the main idea is, Jordan hesitates: "It's about fossil fuels and agreements? And the temperature going up." The first exchange with the chatbot supported Jordan's understanding; the second overloaded it.

INDEPENDENT PRACTICE

PRIYA'S SCREEN

Priya

We have just looked at the difference between the natural and enhanced greenhouse effect. Now choose one cause of the enhanced greenhouse effect and explain how it contributes to the problem.

AI chatbot

One major cause of the **enhanced greenhouse effect** is the burning of **fossil fuels** such as coal, oil, and natural gas.

When fossil fuels are burned to produce electricity, power vehicles, or run industries, they release large amounts of **carbon dioxide (CO₂)** into the atmosphere. Carbon dioxide is a greenhouse gas that traps heat from the Sun. As more CO₂ builds up in the atmosphere, more heat is trapped, causing the Earth's average temperature to rise.

This extra warming strengthens the natural greenhouse effect, leading to problems such as **climate change, rising sea levels, more extreme weather events, and changes to ecosystems**.

In summary, burning fossil fuels contributes to the enhanced greenhouse effect by increasing the amount of carbon dioxide in the atmosphere, which traps more heat and warms the planet.

Priya

Why does CO₂ trap the heat from the sun?

AI chatbot

Carbon dioxide (CO₂) traps heat because its molecules can absorb and re-emit infrared radiation, which is the type of energy the Earth releases after being warmed by the Sun. When sunlight reaches Earth, it warms the surface, which then gives off heat as infrared radiation. CO₂ molecules absorb this heat energy, causing them to vibrate, and then release the energy again in all directions. Some of this re-emitted heat is sent back toward the Earth's surface instead of escaping into space, which keeps the planet warmer. This ability comes from CO₂'s molecular structure, which allows it to interact with infrared radiation in ways that common gases like oxygen and nitrogen cannot. [un +1](#) [modern-physics](#)

Priya

So if we reduce CO₂ emissions would the effect reverse?
How long would that take?

AI chatbot

Reducing CO₂ emissions would help slow and eventually stop further warming, but it would not quickly reverse the enhanced greenhouse effect. This is because carbon dioxide stays in the atmosphere for a very long time, so the extra heat that has already been trapped continues to affect the planet. Even if emissions were reduced to zero, global temperatures would likely stop rising within a few decades, but they would remain higher than normal for centuries. Over much longer timescales, possibly hundreds to thousands of years, temperatures may gradually decrease as CO₂ is slowly removed from the atmosphere by natural processes such as absorption by oceans. [oggardenonline +1](#)



SCIENCE OF LEARNING BUILDING KNOWLEDGE VS COGNITIVE OVERLOAD

Priya is building knowledge step-by-step – each question stems from the previous answer, moving from the distinction to the mechanism to a causal question about reversibility. Each new piece of understanding connects to and deepens what came before. This is how lasting schemas develop: gradually, with new information attaching to existing understanding rather than arriving without context. Her attention is focused on seeking and receiving specific information, without becoming overloaded.

Jordan, by contrast, has drifted back to the pattern Ms Orlando modelled against. His second prompt returns a response covering causes, effects, policy, and science simultaneously – more new information than working memory and weakly established schemas can meaningfully process at once. The chatbot's response is not wrong. It is simply more than Jordan can absorb right now. He is reading, but the conditions for learning are no longer in place. It reveals something more persistent: knowing how to engage productively in one moment does not automatically carry forward to the next. Sustaining focused, deliberate thinking requires students to monitor their own approach as they go – noticing when they are genuinely building understanding and when they are simply accumulating information. This is a skill that needs to be taught, not assumed. It will be explored in more depth in Key Area 4.



MS ORLANDO STEPS IN

Ms Orlando returns to Jordan. She does not address the content of the chatbot's response. She addresses the approach.

MS ORLANDO

"Close that for a moment. Tell me, in your own words, what you understood from your first question."

Jordan explains the natural versus enhanced distinction clearly. Ms Orlando says: "You've got that. Now, what's one thing from that response you actually want to understand better that will help you meet the success criteria?" She points to where it is displayed to focus attention. Jordan thinks. "I guess... why burning fossil fuels makes such a big difference compared to other things." Ms Orlando says: "Ask that. Just that."

JORDAN'S SCREEN – REVISED PROMPT

Jordan

Why does burning fossil fuels produce so much more CO₂ than other activities?

AI chatbot

Burning fossil fuels produces much more carbon dioxide than other activities because fossil fuels are made mostly of carbon, and when they are burned this carbon reacts with oxygen in the air to form CO₂. In addition, fossil fuels contain carbon that has been stored underground for millions of years, so burning them releases large amounts of "new" carbon into the atmosphere very quickly, rather than simply recycling carbon already in the natural carbon cycle. Humans also burn fossil fuels on a very large scale for electricity, transport, and industry, which greatly increases total emissions. Finally, the chemical process of combustion adds oxygen from the air to the carbon in the fuel, meaning that even a small amount of fuel can produce a much larger mass of CO₂.

SCIENCE OF LEARNING RETRIEVAL AND REFOCUSING ATTENTION



Ms Orlando's move here is precise. Before asking Jordan to engage with more content, she asks him to retrieve what he already understood – closing the screen and saying it aloud. This does two things: it consolidates the understanding Jordan built in his first exchange, and it re-establishes Jordan as the active thinker rather than a passive reader of generated text. The follow-up question Jordan generates himself – about fossil fuels specifically – is focused enough that the chatbot's response can connect meaningfully to what he already knows. The extraneous load has been reduced not by simplifying the content, but by restructuring how Jordan is approaching it and focusing his attention back to the success criteria and learning intention.

FIFTEEN MINUTES LATER

Ms Orlando pauses the class. She doesn't share anyone's screen – she puts a question to the room:

"Stop what you're doing. Without looking at the chatbot or your notes, write down in your own words, using full sentences, the most important thing you've understood so far. You have two minutes."

Students close screens or look away before writing. After two minutes, Ms Orlando randomly asks three students to share. A short discussion follows. The chatbot sits idle for the duration.

Before students reopen their screens, Ms Orlando asks two more questions:

"The chatbot has told you things about the enhanced greenhouse effect. How confident are you that what it told you is accurate? Hold up the number of fingers to demonstrate your confidence: 5 for 'very' and none for 'not at all confident.'" [Some teachers ask students to close their eyes while doing this to reduce social desirability bias.]

Ms Orlando surveys the room to get a check on their self-assessed levels of confidence.

"How would you check for the chatbot's accuracy?"

A few students suggest the textbook, the Bureau of Meteorology, their science notes. Ms Orlando doesn't verify each suggestion – she notes them and moves on. But the question has been planted: the chatbot is a starting point, not a final source.

SCIENCE OF LEARNING RETRIEVAL PRACTICE AND DURABLE LEARNING



This is a deliberate retrieval break – a moment of productive effort that interrupts passive reading of chatbot responses. Writing without looking at the screen forces students to retrieve from memory rather than simply recognise information already in front of them. As established earlier in this guide, this kind of effort, even when it feels harder, strengthens long-term retention in ways that re-reading does not. The short discussion that follows gives Ms Orlando real-time insight into what the class has genuinely understood, what remains fragile, and what needs to be addressed before the lesson moves on. The chatbot did not create this moment. The teacher did.

DIGITAL LITERACY EVALUATING AI-GENERATED CONTENT



Before students reopen their screens, Ms Orlando asks them to consider how they would verify what the chatbot has told them. This is a deliberate information literacy move. GenAI produces fluent, confident responses – but it can be factually incorrect, and it cannot verify the sources it draws on. Students who do not develop the habit of checking AI output against reliable sources are not just at risk of acquiring wrong information. They are at risk of not knowing they have done so. Making source evaluation an explicit part of every AI-assisted task – not an afterthought – models critical evaluation of digital content, a skill that develops across all year levels. Ms Orlando does not check the sources for the class. She makes the question a classroom habit. That is where the skill is built.

2. BUILDING KNOWLEDGE AND STRENGTHENING LONG-TERM MEMORY WITH GENAI

Managing attention and cognitive load, as the previous section showed, only gets information into working memory.

That is where the learning starts, but it is not where it stays. For new information to become long-term, organised knowledge, students have to retrieve it, connect it, apply it, and return to it over time (Bartlett, 1932; Anderson et al., 1977; Sweller, 2011). Knowledge is built, not received.

This is where GenAI potentially cuts both ways. Used well, it may push students to process knowledge actively, prompting them to explain, elaborate, and recall rather than simply read. Used poorly, it hands them fluent answers that feel like understanding while quietly bypassing the effortful work that builds it.

HOW KNOWLEDGE BECOMES DURABLE

Learning depends on encoding information into long-term memory and integrating it into existing schemas – the organised knowledge networks that let students interpret new information, make connections, and think efficiently. Think of a schema like a Lego model: when a student has genuinely built one, retrieving it gives them not just the structure but all the relationships between its parts. The key word is built. Schemas are constructed through active processing, not acquired through exposure. Students build stronger, more transferable knowledge when they explain ideas in their own words, generate their own examples, recall knowledge from memory, compare concepts, and connect new learning to what they already know (Craik & Lockhart, 1972; Roediger & Karpicke, 2006). These are not simply good study habits. They are the processes through which knowledge becomes durable – and classroom experiences that bypass them, through passive reading, copying, or surface engagement, tend not to translate into learning.

Two findings are likely to matter most for how teachers use GenAI:

- **Retrieval practice:** recalling knowledge from memory strengthens it far more than re-reading or reorganising it, and the effort consolidates memory in ways that simply seeing the information again does not (Roediger & Karpicke, 2006).
- **Spacing and interleaving:** revisiting knowledge across time and mixing related problems rather than practising them in blocks, improves both retention and transfer, because students must repeatedly retrieve and distinguish between ideas rather than rely on short-term familiarity (Bjork & Bjork, 2011; Cepeda et al., 2006; Kornell & Bjork, 2008).

Both findings point the same way.

WHAT THIS MEANS FOR GENAI IN THE CLASSROOM

Unlike earlier technologies, which displayed fixed content, GenAI responds to what a student says – adapting an explanation, pushing back on a claim, asking a follow-up question. This makes elaborative dialogue possible: a student explains their understanding, the chatbot identifies a gap, the student revises, and the exchange continues. Used this way, GenAI might work less like a textbook and more like a thinking partner.

The risk is the mirror image: because GenAI generates fluent explanations instantly, students can receive a complete, well-organised answer without doing any of the cognitive work that makes new information stick and fit within existing schemas. A clear AI summary feels productive and can create a genuine impression of understanding, without the retrieval, elaboration, or connection-making that actually builds durable knowledge.

GenAI is therefore likely to be most effective when it requires students to do something with knowledge rather than simply receive it: to generate their own retrieval questions and attempt to answer them before checking; to explain a concept and ask the chatbot to

find the gaps; to build and refine analogies through dialogue; to compare related concepts in their own words; or to use feedback that prompts revision rather than confirms an answer.

One well-evidenced application is the tailored worked example. Novice learners acquire knowledge more efficiently by studying worked examples before attempting problems independently, because a worked example reduces problem-solving load while focusing attention on the structure of the solution (Sweller, 2006; Renkl, 2014). GenAI can generate these on demand, adapt them to the concept a student is struggling with, and fade the support as competence develops – from complete examples to partially completed problems, to independent practice. This fading of guidance mirrors what expert teachers do intuitively through gradual release of responsibility models.

The following vignette illustrates how a teacher uses GenAI to strengthen retrieval practice, support elaboration, and consolidate knowledge, while keeping students responsible for the thinking that learning requires.

DURABLE LEARNING IS NOT THE PRODUCT OF A SINGLE CLEAR EXPOSURE, HOWEVER WELL EXPLAINED. IT IS THE PRODUCT OF EFFORTFUL, REPEATED ENGAGEMENT OVER TIME - EXACTLY THE KIND OF EFFORT GENAI CAN POTENTIALLY EITHER PROTECT OR QUIETLY REMOVE.

CLASSROOM EXAMPLE: YEAR 10 PERSONAL DEVELOPMENT, HEALTH AND PHYSICAL EDUCATION

FROM LISTING TO UNDERSTANDING

THE SETUP

Mr Chen's Year 10 Personal Development, Health and Physical Education (PDHPE) class finished an introductory lesson on the social determinants of health four days ago. Students can name the determinants – individual, sociocultural, socioeconomic, and environmental factors – but Mr Chen is not convinced they understand how these factors interact in real situations. Naming a list is not the same as understanding a system.

Today's lesson is designed around that gap. Before students open their chatbots, Mr Chen introduces the success criteria.

By the end of the lesson, students will be able to:

1. clearly explain more than one determinant of health
2. show how they interact in a given example
3. use a chatbot conversation to show your thinking, not just your memory
4. improve your explanation of determinants from the start to the end of the lesson

MR CHEN

"Today I want to find out not just what you know, but how well you can use it. The determinants of health don't work in isolation – they interact. Your job today is to show the chatbot exactly how."

SCIENCE OF LEARNING THE DIFFERENCE BETWEEN RECALL AND UNDERSTANDING



Students who can list the determinants of health may have stored that information in long-term memory – but listing is not the same as understanding. Genuine schema development requires knowledge to be organised into connected structures where relationships between ideas are as important as the ideas themselves.

Mr Chen's lesson is designed to surface exactly this gap: the difference between students who know what the determinants are and students who understand how they work together.

STAGE ONE - RETRIEVAL WITHOUT THE CHATBOT

Laptops stay closed. Mr Chen gives students five minutes to write everything they can recall about the determinants of health without notes, laptops, phones, or the chatbot.

Retrieval Task – Individual, No resources

Write down everything you know about the determinants of health. Include what they are, what they mean, and – if you can – how they might connect to each other. You have five minutes.

After five minutes, Mr Chen asks students to read back what they wrote and circle anything they feel uncertain about. Then he puts a scenario on the board.

The Scenario

Anika is 16 and lives in a regional town three hours from the nearest city. Her family rents a house in an area with limited public transport. She works part-time after school to help with household expenses. She has been feeling persistently tired and is struggling to concentrate at school but hasn't seen a doctor in two years.

MR CHEN

"Don't open the chatbot yet. Using only what you already know, write down which determinants of health are affecting Anika – and explain how they are connected to each other in her situation. Three minutes."

SCIENCE OF LEARNING RETRIEVAL PRACTICE AND SPACED LEARNING



Mr Chen has structured the opening of the lesson around two retrieval attempts before the chatbot is opened. The first – free recall of the determinants – activates existing schemas and surfaces gaps. The second – applying those determinants to Anika's scenario – requires something more demanding: retrieve, select, and connect knowledge under the pressure of a real context. This is the kind of effortful retrieval that builds durable, transferable understanding. The four-day gap since the initial lesson is not accidental – spacing retrieval across time strengthens memory significantly more than immediate review.

STAGE TWO - EXPLAIN AND BE CHALLENGED

Students open their chatbots. Mr Chen gives them a structured task: explain in full sentences how the determinants of health are affecting Anika, then ask the chatbot what they are missing or oversimplifying.

Explicit Instruction – Provided to all students

Explain to the chatbot how the determinants of health are affecting Anika in the scenario.
Be as specific as you can. Then ask: “What am I missing or oversimplifying in this explanation?”

Two very different exchanges emerge.

JOHN'S SCREEN

John

Anika is affected by socioeconomic determinants because her family doesn't have much money and environmental determinants because she lives far from the city. These are affecting her health because she can't access healthcare easily. What am I missing or oversimplifying in this explanation?

AI chatbot

You've identified two relevant determinants and made a connection to healthcare access – that's a reasonable start. A few things worth pushing further: You've named socioeconomic and environmental factors, but Anika's situation also involves individual factors – her fatigue and concentration difficulties could reflect her health status, sleep, nutrition, or stress levels, which are themselves shaped by her circumstances. There's also a sociocultural dimension worth considering: what role might community, family expectations, or cultural attitudes toward help-seeking play in why she hasn't seen a doctor in two years? And here's a bigger question: you've listed the determinants separately, but how are they interacting with each other in Anika's situation? Does living far from the city make the financial pressure worse, or does the financial pressure make the distance more of a barrier?

John

I think they make each other worse like if she had money she could afford transport to get to a doctor but because she doesn't have money and lives far away both things together mean she just doesn't go.

AI chatbot

That's a really important insight – what you're describing is sometimes called the 'compounding' or 'accumulation' effect of health determinants. When multiple determinants work against a person simultaneously, the combined effect on health is often greater than any single factor alone. This is central to understanding health inequity – it's not just that some people face one disadvantage, but that disadvantages tend to cluster together. Can you think of another example from Anika's scenario where two determinants are compounding each other?

SCIENCE OF LEARNING ELABORATIVE INTERROGATION AND CONNECTED SCHEMAS



John began with a list – two determinants, one consequence. The chatbot's challenge did not correct him; it pushed him to move from listing to connecting. The insight John arrives at – that the determinants compound each other – is not information he received. It is a relationship he constructed under the pressure of a question that required him to think beyond the list. This is elaborative interrogation in action: being pushed to explain how and why, rather than what, forces the kind of deep processing that builds connected, organised schemas rather than isolated facts.

ZARA'S SCREEN

Zara

Anika is affected by multiple interacting determinants. Socioeconomically, her family's financial situation means she works part-time, which reduces her time for rest and probably affects her nutrition and sleep which contributes to her fatigue. Environmentally, she lives in a regional area with limited transport, which creates a physical barrier to healthcare access. These two factors compound each other: even if she wanted to see a doctor, the distance combined with the cost of transport and potentially taking time off work makes it very unlikely she will. The sociocultural factor might also be relevant as in some communities there are norms around not seeking help unless things are serious. What am I missing or oversimplifying in this explanation?

AI chatbot

That's a well-structured analysis that identifies both individual determinants and their interactions – the compounding point is particularly strong. Two things worth developing: First, you've focused on barriers to accessing healthcare, but the determinants are also shaping Anika's health status independently of healthcare access. Her fatigue and concentration difficulties might be direct outcomes of poor sleep, inadequate nutrition, or chronic stress – which are themselves consequences of the socioeconomic pressures, not just symptoms waiting for a doctor. Second, there's a concept in health promotion called the 'social gradient of health' – the idea that health outcomes aren't just about poverty versus wealth, but that there's a continuous relationship between social position and health across the whole population. Where does Anika sit on that gradient, and what does that suggest about the kind of interventions that would actually make a difference for someone in her situation?

Zara

I hadn't thought about the health effects being separate from healthcare access like the stress and fatigue aren't just because she can't see a doctor, they're caused by her actual life circumstances. So even if she could access a doctor easily, she'd still be experiencing those health effects. That means fixing healthcare access alone wouldn't fix the problem.

AI chatbot

Exactly and that's the insight that underlies most of the debate about how to reduce health inequity. If the determinants themselves are causing the poor health outcomes, then healthcare is treating symptoms rather than causes. This is why public health approaches focus on the social determinants – income, housing, education, employment – rather than just building more hospitals. Does this change how you'd think about what an effective health promotion strategy for someone like Anika would look like?

SCIENCE OF LEARNING KNOWLEDGE AT THE EDGE OF COMPETENCE



Zara's initial explanation was already sophisticated – she identified compounding effects and sociocultural factors without prompting. The chatbot's role here is not gap-filling but extension: introducing the social gradient concept and pushing Zara toward a structural understanding of health inequity that goes beyond Anika's individual circumstances. The insight Zara arrives at – that fixing healthcare access wouldn't fix the underlying problem – is a genuinely new understanding, not a retrieval of something she already knew. This is schema elaboration at the edge of competence: new knowledge connecting to and reorganising existing understanding rather than simply adding to it.

MR CHEN PAUSES THE CLASS

After twelve minutes, Mr Chen stops the class. He does not share anyone's screen.

MR CHEN

"Close the chatbot. In your notebook, write one sentence that explains how the determinants of health interact in Anika's situation – not a list, one sentence that shows the connection. Then write one question you now have that you didn't have at the start of the lesson. And one more thing – the chatbot introduced some concepts you may not have seen before, like the social gradient of health. Before next lesson, I want you to find one other source that either confirms or adds to what it told you."

Students write in silence. Mr Chen moves around the room reading. He is looking for two things: whether students are still listing or have shifted to connecting, and what kinds of new questions are emerging. After three minutes, he asks four students to share their questions. He writes them on the board without answering them. They will form the basis of next week's lesson on health promotion.

SCIENCE OF LEARNING CONSOLIDATION, RETRIEVAL, AND GENERATIVE THINKING



The one-sentence task is harder than it looks. Writing a sentence that shows connection, rather than listing, requires students to have built a relational understanding, not just accumulated information. Students who are still thinking in lists will find this difficult, and that difficulty is informative: it tells Mr Chen, and the student, that the schema is not yet organised around relationships. The new question prompt is equally deliberate: a student who can ask a better question than they started with has built new knowledge. Mr Chen's decision to write those questions on the board and defer them to next lesson is a spaced learning decision – the questions will be retrieved again in a week's time, with the interval doing its consolidating work in between.

DIGITAL LITERACY EVALUATING AI-GENERATED CONTENT



The chatbot in this exchange introduces concepts neither John nor Zara have encountered before – the compounding effect of health determinants and the social gradient of health. Both are real and accurately described here, but students have no way of knowing that without checking. Mr Chen's instruction to verify these concepts against another source before next lesson is a deliberate habit-building move: AI output requires the same critical evaluation as any other source. That habit is most effectively built not as a separate digital literacy lesson but embedded in the moment when students have just encountered something new and fluently explained – which is precisely when the need for verification is least obvious and most necessary.

STAGE THREE - REBUILD THE EXPLANATION

Mr Chen asks students to return to their original retrieval attempt – the first thing they wrote before the chatbot – and compare it with what they understand now.

MR CHEN

“Read what you wrote at the start. Then, without using the chatbot, write a new explanation of how the determinants affect Anika. Make sure you:

- write in full sentences
- name more than one determinant
- explain how they interact, not just list them
- show what you understand now that you didn’t before”

John’s two attempts sit side by side in his workbook.

John’s Workbook – Beginning of Lesson

Anika is affected by socioeconomic and environmental determinants. She doesn’t have money and lives far away so she can’t get to a doctor.

John’s Workbook – End of Lesson

Anika’s health is affected by multiple determinants that compound each other. Her family’s financial situation means she works part-time, leaving less time for sleep and rest, which directly contributes to her fatigue. The environmental barrier of living in a regional area with limited transport combines with the financial barrier to make healthcare access very unlikely – not just difficult but effectively out of reach. These determinants don’t work separately: the financial pressure makes the distance a bigger problem, and the distance makes the financial pressure harder to manage. Even if healthcare access improved, the underlying socioeconomic conditions would still be causing harm to her health.

SCIENCE OF LEARNING VISIBLE SCHEMA DEVELOPMENT



The distance between John’s two attempts is a visible record of schema development. The second explanation is not longer because John has read more – it is more organised, more relational, and more analytically precise because he has had to retrieve, elaborate, be challenged, and reconstruct. The shift from listing to connecting is not cosmetic: it reflects a genuine change in how the knowledge is organised in long-term memory. Schemas built around relationships – how things interact – are significantly more transferable than schemas built around lists of facts. When John encounters a different scenario next week, the relational framework he has built today will allow him to apply his understanding in ways that simple recall never could.

3. APPLICATION AND TRANSFER: USING KNOWLEDGE IN NEW SITUATIONS

The first two key areas dealt with how students take in information and build it into lasting schemas: first by managing attention and cognitive load, then by consolidating what is learned into organised, durable memory. Key Area 3 asks what happens when the situation changes: can students use what they know in a context they have not seen before?

This is the question of transfer – the ability to apply knowledge flexibly to new problems, unfamiliar contexts, and situations that look nothing like the ones in which the learning first happened. Transfer is the goal of most teaching. It is also, as the science of learning makes clear, one of the hardest things to achieve and is often assessed to test depth of understanding, such as in summative examinations.

WHY TRANSFER IS HARDER THAN IT LOOKS

Students often learn knowledge in one context – a specific example, a worked problem, a familiar scenario – and then fail to recognise that the same principle applies when the surface details change. A student who can solve a force problem involving a car on a flat road may not see that the same principles apply to a boat on water. This is not laziness or inattention; it reflects something important about how novice learners store and retrieve knowledge.

Research consistently shows that novices tend to focus on the surface features of a problem – what it looks like – rather than its underlying structure – what kind of problem it is (Chi et al., 1981). Expert problem solvers do the opposite. They recognise the deep structural pattern quickly, retrieve the right approach from long-term memory, and apply it with little conscious effort. The difference is not intelligence; it is the depth and organisation of the schemas built through deliberate practice and experience. It is for this reason that problem-solving is not a generic skill that can be taught – it requires background knowledge, stored as schemas, from which abstractions can be made over time.

This is why transfer is so closely tied to schema development. Applying knowledge to a new problem requires schemas organised around underlying principles rather than

surface features – around why something works, not just what to do. A student who has learned a procedure without understanding the principle behind it has no reliable way to recognise when that principle applies somewhere new. The surface has changed, and that is all they can see.

This capacity also develops across adolescence. Younger secondary students may have genuine knowledge in a domain and still struggle to see past surface features to underlying structure – not because they haven't learned enough, but because the ability to think abstractly about what makes problems similar grows with experience and maturity. Transfer needs to be explicitly taught at every stage: through varied examples, through tasks that ask students to explain why an approach works, and through practice applying the same principle across different contexts.

For students who are still developing their knowledge in this area, working through an unfamiliar problem is genuinely effortful. Without well-developed schemas to lean on, students have to hold several things in working memory at once: what they know, what they need to find out, and how to get from one to the other. This is one of the elements that makes novel problems so demanding, and helps explain why students so often reach for the first strategy that comes to mind rather than the most appropriate one.

WHAT HELPS STUDENTS TRANSFER KNOWLEDGE

The science of learning points to several conditions that make knowledge transfer more likely:

- **A strong knowledge base** – students cannot transfer what they do not know. Shallow familiarity with a topic does not produce flexible understanding; depth does.
- **Varied practice** – encountering the same underlying principle across different contexts and problem types loosens knowledge from any one surface form and makes it more available for use in new ones (Bjork & Bjork, 2011).
- **Self-explanation** – students who can explain why a solution works, not just what the answer is, develop more transferable understanding than those who simply reach the right answer (Chi et al., 1989).

The last condition is where GenAI potentially becomes most relevant. Making reasoning explicit – slowing down to say “this works because...” – is what helps knowledge become flexible enough to transfer. It is effortful. It is also exactly the kind of productive cognitive work that builds lasting understanding.



WHAT THIS MEANS FOR GENAI IN THE CLASSROOM

GenAI introduces a specific tension here. It can solve problems quickly, accurately, and with clear working shown, which genuinely supports learning when students use it to check their reasoning, find where their thinking went wrong, or understand a principle they have misapplied. But when students use it to generate answers without first attempting the problems themselves, they skip the effortful retrieval and problem solving that builds the very schemas transfer depends on.

The risk is not simply that students get answers without working; it is that a fluent, well-structured solution can give the impression of understanding without producing it.

GenAI is potentially most useful for transfer when it acts as a thinking partner rather than a solution provider – asking students to explain their reasoning, name the underlying principle in a problem, or apply a method to a different scenario. Used this way, it helps students move from “I know how to solve this type of problem” to “I understand why this approach works”, which is the shift that makes knowledge transferable. The distinction worth holding onto is precise. Using GenAI to retrieve

information – looking up a fact, checking a definition, confirming a calculation – is often entirely appropriate, and is cognitively no different from consulting a textbook or a search engine. The concern is using it to generate the thinking, reasoning, or construction the student should be doing. The line is not between using AI and not using it, but between using it to extend thinking and using it to replace it.

This has direct implications for senior secondary students preparing for closed-book examinations.

An unseen exam question is a transfer task: it requires students to recognise an underlying principle and apply it to a context they have not encountered before.

That capacity is built through the effortful problem-solving GenAI can so easily shortcut. Students who have practised retrieving and applying knowledge independently are practising exactly the skill the exam requires. Students who have relied on AI to generate solutions have practised something else entirely – and may not discover the difference until the exam room makes it unavoidable. The following classroom example illustrates what this looks like in practice.

A STUDENT WHO READS A CLEAR AI-GENERATED WORKED SOLUTION MAY FEEL THEY UNDERSTAND - AND GENUINELY BELIEVE THEY DO - WITHOUT HAVING DONE THE COGNITIVE WORK THAT WOULD CONSOLIDATE THAT NEW LEARNING IN LONG-TERM MEMORY AND MAKE IT AVAILABLE IN A NEW CONTEXT.



CLASSROOM EXAMPLE: YEAR 10 ENGLISH

SEEING PAST THE SURFACE

THE SETUP

Ms Clifford's Year 10 English class has spent the past two weeks analysing persuasive texts. They have worked through a series of opinion pieces and political speeches, building a shared vocabulary for identifying how writers construct arguments: appeals to authority, emotional appeals, rhetorical questions, inclusive language, repetition, and the strategic use of evidence. Students can identify these techniques confidently in the texts they have studied.

Today's lesson asks them to do something harder. Ms Clifford puts a social media post on the board – a screenshot from a popular health and wellness influencer. It is visually polished, informally written, and built around a personal story. It looks nothing like the political speeches and opinion pieces students have been analysing. Ms Clifford says nothing about persuasive techniques. She just asks one question.

MS CLIFFORD

"Is this text trying to persuade you? How do you know?"

Students discuss briefly in pairs. Most agree it is persuasive. Several point to the product being promoted. Ms Clifford nods and moves on.

MS CLIFFORD

"Now the harder question. Use the same analytical tools you've been building over the past two weeks to analyse how this text works. Not what it's selling – how it persuades."

Students open their chatbots. Ms Clifford gives the class a starting prompt and asks them to work through it before moving to their own questions.

STARTING PROMPT FOR CHATBOT – PROVIDED TO ALL STUDENTS

I'm going to analyse how a social media post persuades its audience, using the persuasive techniques we've studied in political speeches and opinion pieces – appeals to authority, emotional appeals, rhetorical questions, inclusive language, repetition, and the strategic use of evidence. I'll share my analysis in my next message.

Teacher Instruction

Explain to the chatbot how you're using the persuasive techniques we've studied to analyse the post. Remember to analyse how each technique works, not just name it. Then ask: "What am I missing, and is my analysis actually explaining the technique or just naming it?"

DIGITAL LITERACY READING CONTENT STUDENTS CANNOT YET POST



Under the social media age restrictions that took effect in December 2025, students under 16 cannot create or maintain accounts on platforms such as Instagram, TikTok, or YouTube. The restriction applies to accounts, not to seeing the content: students still encounter this material on shared screens, in screenshots, and once they turn 16. Ms Clifford displays the post herself, so no student needs an account to take part. Teaching students to see how this kind of content persuades is not at odds with the age restrictions – it serves the same purpose, preparing young people to meet social media as critical readers rather than passive ones.

RUBY'S FIRST EXCHANGE

Ruby is a capable student. She identified persuasive techniques accurately in every text studied so far. Ms Clifford watches her screen as she types her first response to the chatbot.

RUBY'S SCREEN

Ruby

Here is my analysis. The influencer uses an emotional appeal by sharing a personal story about her health journey. She uses inclusive language like 'we all know how it feels' to connect with the audience. She also uses repetition of the word 'natural' throughout the post. What am I missing, and is my analysis actually explaining the technique or just naming it?

AI chatbot

You've correctly identified three techniques and named them accurately. But the second part of your question is worth sitting with – are you explaining them or naming them? Let's take the emotional appeal. You've said she shares a personal story about her health journey. What does that actually do to the reader? Why is a personal story more persuasive than, say, a statistic or an expert quote in this context? And the repetition of 'natural' – what does that word mean to the audience she's targeting, and why does repeating it matter? Naming a technique tells you what's there. Explaining it tells you how and why it works. Which of your three analyses gets closest to doing that second thing?

Ruby

I think the inclusive language one is closest because I said it connects with the audience – but I didn't explain why that connection makes them more likely to trust her or buy the product.

AI chatbot

That's a precise observation. So what is the connection between feeling understood and being persuaded? Why does an audience who feels seen by a writer become more receptive to what that writer is recommending?

Ruby

Because if you feel like someone gets your experience, you trust them more. And if you trust them, you're more likely to believe what they say about the product.

AI chatbot

Exactly – and that's the underlying mechanism. The technique works because it builds credibility through identification rather than expertise. Now here's a question worth thinking about: is that the same mechanism as anything you've analysed in the political speeches or opinion pieces? Does it remind you of any technique that worked differently on the surface but achieved the same thing underneath?



SCIENCE OF LEARNING NAMING VERSUS EXPLAINING

Ruby's first response is accurate but thin. She has identified the techniques correctly – but identification alone does not demonstrate understanding. The chatbot's challenge has pushed her toward the harder cognitive task: explaining the mechanism, not just the label. This is where transfer begins. Understanding why a technique works, not just what it is called, is what allows a student to recognise it when it appears in a completely different form. Ruby cannot yet see the connection to the political speeches because she has been thinking about techniques as categories rather than as mechanisms. The chatbot is pushing her toward the level of abstraction that makes transfer possible.



DIGITAL LITERACY USING PRIOR KNOWLEDGE TO EVALUATE AI RESPONSES

Ruby has two weeks of analytical work to draw on. Before accepting the chatbot's challenge or its framing of the techniques, she can ask: does this fit with what I already know about how persuasion works? That question — does this response make sense considering what we've already learned? — is a more sophisticated information literacy move than source-checking alone. It uses schema depth as the verification mechanism. Students with enough prior knowledge can evaluate AI responses from the inside, not just by consulting an external source. Ms Clifford's deliberate sequencing — two weeks of analytical work before the chatbot is introduced — is what makes that evaluation possible.

THE TRANSFER MOMENT

Ruby pauses. Ms Clifford, passing behind her, sees her staring at the screen rather than typing. She stops.

MS CLIFFORD

"What are you thinking?"

Ruby: "In the political speech – Whitlam – he kept saying 'it's time.' That's repetition. But it's also kind of the same thing as the inclusive language here, isn't it? Like, both of them are making you feel like you're part of something. Like you belong to this moment or this group."

Ms Clifford does not answer immediately. She asks: "Is that the same technique, or the same effect achieved by different techniques?"

Ruby thinks. "Same effect. Different technique. Because 'it's time' is repetition and the 'we all know' thing is inclusive language. But they're both trying to make you feel like you're already on board."

MS CLIFFORD

"Write that down before you do anything else. Exactly what you just said."

SCIENCE OF LEARNING THE TRANSFER MOMENT



What Ruby has just done is the cognitive event that Key Area 3 is about. She has seen past the surface features – a 1970s political speech and a contemporary social media post look nothing alike – to the underlying persuasive mechanism both texts are using: creating a sense of belonging and shared identity to lower the audience's resistance. This is transfer. This wasn't the result of Ruby being told to look for connections. It happened because the chatbot's questions pushed her to understand why techniques work, and that deeper understanding made the structural similarity visible. Ms Clifford's instruction to write it down immediately is deliberate – the insight is fragile at this stage and articulating it in writing begins to consolidate it as transferable knowledge rather than a passing observation.

DEEPENING THE ANALYSIS

Ruby returns to the chatbot with what she has noticed.

RUBY'S SCREEN

Ruby

I've noticed the inclusive language in the social media post and the repetition of 'it's time' in the Whitlam speech are different techniques, but they're achieving the same effect – making the audience feel like they already belong to a group that agrees with the speaker. Is that a principle I can apply to other texts?

AI chatbot

Yes and you've just named it more precisely than most analytical frameworks do. What you're describing is sometimes called identification or affiliation – the idea that persuasion works not only by presenting evidence or argument, but by making the audience feel they share an identity with the speaker or with others who hold the same view. Once you belong, disagreeing feels like a betrayal of the group. This mechanism appears across very different text types – speeches, advertisements, social media posts, political manifestos – with completely different surface features. Now that you can see the mechanism, try this: go back to your original analysis of the emotional appeal – the personal health story. Does the same principle apply there? Is the personal story also working through identification?

Ruby

Yes she's not just saying 'this product helped me.' She's saying, 'I was where you are, and now I'm here.' She's positioning herself as someone who used to be part of the audience's group, so they identify with her past self and trust her current recommendation.

AI chatbot

That's a sophisticated reading. And notice what you've done – you've used the same underlying principle to explain three different surface features of the one text. That's what analytical frameworks are for: not to label things, but to reveal how they work. Could you now write a thesis sentence that makes an argument about how this text persuades, using that principle as your central claim?

SCIENCE OF LEARNING SELF-EXPLANATION AND TRANSFERABLE KNOWLEDGE



Ruby's analysis has moved through three distinct stages. She began by naming techniques – accurate but surface-level. The chatbot pushed her to explain mechanisms – why techniques work, not just what they are. And now she has abstracted a principle – identification and affiliation – that applies across multiple techniques within the one text and across completely different text types. This is the self-explanation effect in action. By being required to articulate why, repeatedly and precisely, Ruby has built an understanding of persuasion that is no longer tied to the specific texts she has studied. She can take this principle into an exam, a new text, or a different medium and apply it with genuine understanding rather than pattern matching.

MS CLIFFORD BRINGS THE CLASS TOGETHER

Near the end of the lesson, Ms Clifford asks students to close their chatbots. She puts one question on the board.

MS CLIFFORD

"Without looking at anything – write one sentence that explains the most important thing you now understand about how persuasion works. Not a technique. A principle."

Students write in silence. Ms Clifford reads several aloud (anonymously) and the class discusses. Ruby's sentence is one of them:

Persuasion works not just by presenting arguments, but by making the audience feel they already belong to the group that agrees with you.

MS CLIFFORD

"If that's true – where else might you see that principle at work? Beyond texts we've studied in class?"

The discussion that follows ranges across advertising, social media algorithms, and political messaging. Students are applying a principle they extracted from two very different texts to contexts they have never been asked to analyse before.

That is transfer.

SCIENCE OF LEARNING CONSOLIDATION AND FAR TRANSFER



The closing task – writing a principle rather than listing techniques – is the same move Mr Chen used in Key Area 2, but the cognitive demand is higher here. Students are not summarising what they learned. They are abstracting from specific examples to a generalisable claim. The class discussion that follows takes that abstraction further: students applying the principle to contexts entirely outside the curriculum. This is what the science of learning calls transfer – the application of knowledge to situations genuinely remote from the original learning context. It is the hardest kind of transfer to achieve and the most valuable. It does not happen through exposure to more content. It happens when students have been pushed to understand principles deeply enough that they can see them operating across radically different surfaces.

4. METACOGNITION AND SELF-REGULATION: TAKING OWNERSHIP OF LEARNING

The first three key areas have followed a deliberate sequence.

Key Area 1 looked at how students attend to new information and manage cognitive load (attention and working memory). Key Area 2 looked at how knowledge is built into lasting schemas through retrieval and elaboration (long-term memory). Key Area 3 asked whether students can use that knowledge when the context changes (transfer). Each of these depends on the fourth – because none of them happens reliably unless students can monitor their own thinking and take responsibility for directing it. This is the territory of metacognition and self-regulated learning.

OF THE FOUR, THIS CAPABILITY IS THE MOST DIRECTLY AT RISK WHEN GENAI IS USED WITHOUT INTENTION - AND THE ONE MOST WORTH DEVELOPING DELIBERATELY.

WHAT METACOGNITION MEANS AND WHY IT MATTERS

Metacognition refers to an individual's awareness and understanding of their own thinking processes (Flavell, 1979). It has two key components. Metacognitive knowledge is understanding how you learn, including which strategies work, when to apply them, and where your understanding is strong or needs further development. Metacognitive regulation involves using that knowledge to manage your learning by planning an approach, monitoring your progress, and adjusting when something is not working (Brown, 1980; Flavell, 1979).

These two work together. A student who knows they rush through reading tasks but cannot slow down has knowledge without regulation. A student who adjusts instinctively but cannot say why is regulating without awareness. The most effective learners do both – they notice what is happening in their thinking and they act on it.

Self-regulated learning extends this to motivation and behaviour: whether students actually use what they know about their own learning or take the path of least resistance instead (Zimmerman, 2000; Panadero, 2017). Picture two students working on the same maths problem. One notices a recurring error, identifies the underlying misconception, revisits the content,

and practises until it holds. The other pushes through without noticing the error, quietly reinforces it, and ends up discouraged. The difference is not ability. It is awareness, and the willingness to act on it.

Self-regulation does not look the same for every learner. This is partly a matter of knowledge – and partly a matter of development. The capacity to monitor your own thinking, catch errors, and adjust your approach grows across adolescence, partly because the prefrontal regions underpinning executive function and self-regulation continue to develop into early adulthood, although the timing varies between individuals (Blakemore & Choudhury, 2006; Blakemore, 2018). Junior secondary students may possess motivation and reasonable domain knowledge and still find self-regulation difficult – not because they are unwilling but because the metacognitive skills involved are still forming. Teacher-provided structure is therefore not a crutch for weaker students but developmentally appropriate support for all students at earlier stages, which can be gradually withdrawn as capacity grows.

Students who are still building foundational schemas face a compounding difficulty: they often cannot judge their own understanding accurately because

they do not yet know enough to know what they are missing. Metacognitive development extends into adulthood, with sophisticated metaknowing emerging gradually (Kuhn, 2000). Without sufficient domain knowledge, students may lack the reference points needed to evaluate their own understanding accurately. Re-reading produces a feeling of fluency that is easily mistaken for understanding – what Bjork and colleagues call the illusion of knowing (Bjork, Dunlosky & Kornell, 2013). These students benefit from greater external support, including clear task structures, frequent checks for understanding and

explicit feedback. As knowledge and expertise grow, students get better at monitoring their own thinking and adjusting their approach without prompting. That progression has direct implications for how teachers design GenAI tasks across year levels and varying levels of expertise.

These principles are particularly important in senior secondary contexts. In closed book examinations and other high-stakes assessments, students may need to demonstrate their understanding independently, without access to AI tools. The habits promoted in this guide – retrieving from memory, monitoring understanding,

identifying gaps before they become errors, constructing arguments independently – prepare students for these conditions. A student who has used GenAI the way this guide describes, as a tool that requires them to think rather than one that thinks for them, should be more capable of working independently under exam conditions, not less. The risk runs the other way: students who have used GenAI to generate responses rather than build understanding may find, in an exam room, that the fluency they experienced was the tool's, not their own.



THE ROLE OF FEEDBACK

Self-regulation does not develop in a vacuum. Students need feedback to know whether their understanding is accurate, their reasoning is sound, and whether their approach is working. Without feedback, even highly motivated learners lack the information needed to evaluate their thinking effectively, leaving them navigating without a map. This is not only a practical matter – it reflects how the brain learns. The brain is constantly predicting what should happen and correcting itself when those predictions prove wrong, which is why the error signal – the gap between what a student expected and what actually happened – is one of the basic drivers of learning (Dehaene, 2020).

Research on feedback identifies four levels at which it can operate:

- **Task** – whether an answer is right or wrong.
- **Process** – how a student approached the task and what they might do differently.
- **Self-regulation** – how well the student is monitoring and managing their own learning.
- **Self** – evaluative comments about the person rather than the learning (Hattie & Timperley, 2007).

The most powerful feedback works at the process and self-regulation levels – it tells students not just what went wrong but how to think about the problem differently. Self-level feedback – telling a student they are clever or gifted – is the least effective, because it ties outcomes to fixed traits rather than to strategies they can change.

GenAI has genuine potential at the process and self-regulation levels. A well-designed chatbot can respond to a student's reasoning in real time, pinpoint where the thinking breaks down, and push them to reconsider it – at a scale and immediacy no single teacher

with a class of thirty could achieve. When a student explains their understanding and asks where it is incomplete, a well-prompted response makes them evaluate and revise their thinking. That is process feedback, and it is itself a metacognitive act. The design implication for teachers is clear: have students ask the chatbot to identify the step where their reasoning became unclear or incomplete, not to check whether the answer is right.

The risk is that GenAI feedback defaults to the task level – confirming whether an answer is right rather than interrogating how the student got there. That is the least useful kind for developing self-regulation, because it asks for no reflection. A verdict with no interrogation teaches a student almost nothing about how to think differently next time. The task is to design interactions where the verdict is the start, not the end – prompts that make students explain, justify, and reconsider rather than submit and receive.

The feedback process serves another purpose that GenAI cannot replace: it informs teachers. When reading a set of drafts or exit tickets, or listening to students explain their reasoning, a teacher builds a picture of what the class has grasped, where misconceptions are clustering, and what to teach next. That pedagogical knowledge cannot be outsourced. If GenAI absorbs all the feedback students receive during practice, teachers lose the occasions on which that knowledge is built – and with it, decisions only a teacher can make.

This is also where GenAI opens up something new for assessment. A chatbot interaction log is a record of thinking, not just a record of answers.

The goal is not to replace teacher feedback but to complement it by using GenAI for immediate,

practice-level responses, freeing teachers for what only they can do: read their class, exercise professional judgement, and decide where learning goes next. Used well (as in the History example that follows) the exchanges become a source of diagnostic information for the teacher, not a substitute for the teacher's diagnostic eye. The patterns Ms Ferreira reads in her students' agent exchanges tell her something a marked set of essays alone could not: not just what students got wrong, but which distinction they are not yet making. The agent has not replaced her professional judgement. It has sharpened it.

Where a finished piece of work shows only the destination, the log shows the route – the questions a student asked, how they refined their reasoning when challenged, and where a misconception surfaced. In this sense it works like mini-whiteboard, making student thinking visible in real time – but it goes further, capturing how understanding develops step by step. Used this way, the log becomes a rich source of formative assessment: it tells a teacher not just whether a student has understood, but where their reasoning is breaking down and what to teach next.

There is a more ambitious possibility worth flagging. As GenAI makes finished products less reliable as evidence of a student's own thinking, the process captured in an interaction log may become a more trustworthy basis for judging learning than the product itself. Using logs in this way – to assess the quality of a student's questioning, reasoning, and self-correction – is promising but not straightforward: it raises questions of privacy, consistency and what exactly is being judged, and it works only when the task has been designed so that the thinking, not the tool, is doing the work.

**A CHATBOT INTERACTION LOG IS A RECORD
OF THINKING, NOT JUST A RECORD OF ANSWERS.**



THE ROLE OF MOTIVATION

Cognitive strategies alone do not decide whether students learn. A student who understands retrieval practice but never uses it, or who knows they should ask a follow-up question but takes the easy answer, has a motivation problem, not a knowledge one. Self-regulation requires both skill and the will to apply it. What shapes that will matters for any teacher thinking carefully about how GenAI fits into their classroom.

Motivation is strengthened when three psychological needs are met: autonomy, competence and relatedness (Ryan & Deci, 2000). When students feel a sense of self-direction and ownership, experience growth through effort, and feel connected to others, they are more likely to persist through difficulty, seek feedback, and engage deeply enough to build lasting understanding. When they are not – when learning feels controlled, effortless, or disconnected – students drift toward surface strategies, risk avoidance, and disengagement.

GenAI puts each of these three needs at risk. It can reduce **autonomy** by making the student a consumer rather than a producer of the work. It can undermine **competence** by removing the desirable difficulty that comes from working hard at something not yet understood. The student is spared the frustration, but also the satisfaction, and over time they are left less confident in their own ability

to work things out. This confidence is what Bandura (1997) called self-efficacy – the belief in one's own capability to succeed at a task – and its strongest source is mastery experience, actually succeeding at something hard. Each time GenAI removes the struggle, it removes one of those experiences, leaving the student not only less practised but less convinced they could have worked it out themselves. It can also weaken **relatedness** by substituting human interaction with a tool that does not know or care about the student or their success.

For students who are already disengaged or isolated, replacing human interaction with an AI interaction is not a neutral swap – it removes one of the things most likely to keep them in the room.

These risks are not evenly distributed across all students. Research on self-regulated learning consistently finds that students from disadvantaged backgrounds, students with learning differences, and students with limited prior knowledge in a domain tend to have weaker metacognitive habits (Pintrich, 2004; Zimmerman, 2000). We would expect these students to be the most likely to use GenAI in ways that weaken those habits further. The student who already struggles to monitor their understanding is the most exposed to the false confidence a fluent GenAI answer can create. Teachers working in schools with high proportions of students from

disadvantaged communities need to watch that the tool is not widening the very gaps it could help close.

There is a further link between motivation and cognitive load. When a task places needless burden on working memory, students do not only struggle to think; they also lose the willingness to keep trying (Feldon, 2007). But not all difficulty is unnecessary. Reducing the load that comes from poor task design helps students stay engaged, whereas removing the desirable difficulty at the heart of learning removes the very effort that builds understanding. The goal is not to make learning comfortable. It is to make the difficulty worth it.

For teachers, this means designing GenAI tasks is a motivational decision as much as a cognitive one. The question is not only whether a task manages load or supports retrieval but whether it gives students a genuine reason to invest in their own learning. When students own what they produce, feel the satisfaction of working through something hard, and learn alongside teachers and peers who know and care about them, they bring the effort and self-awareness the science of learning identifies as essential. GenAI is most valuable where those conditions already exist – extending human learning rather than substituting for it. Creating those conditions is the teacher's work, and changes in technology do not change that.

WHAT THIS MEANS FOR GENAI

The conditions motivation needs – ownership, genuine effort, human connection – are exactly the conditions that poorly designed GenAI use erodes. GenAI can supply information, generate responses, and ask questions. It cannot care about a student's progress, notice when someone is about to give up, or create the shared endeavour that keeps learners in the room. Using it well means attending not just to what students know, but to how they experience themselves as learners.

The deeper challenge is not simply that students might use it to cheat. It is that, because GenAI produces fluent, confident answers instantly, it can remove the conditions under which metacognitive awareness develops. That awareness is built in moments of difficulty – when a student hits confusion and has to choose whether to push through, look something up, ask for help, or step back and rethink. Those are moments where self-regulation is practised and built. Hand over an instant answer and the student is spared the struggle, and the metacognitive opportunity inside it is lost.

Early GenAI-specific evidence supports this concern, though it is still limited. In a randomised study comparing learners supported by ChatGPT, a human expert, a writing analytics tool, and no support, Fan et al. (2025) found that ChatGPT improved the quality of the work students produced but did not lift their intrinsic motivation or self-regulated learning, and may encourage what the authors call “metacognitive laziness” – the offloading to the tool of the monitoring and evaluation that learning depends on. A large field experiment in high-school mathematics points the same way: Bastani et al. (2025) found that students with access to a standard

chatbot practised more successfully but, once it was removed, scored 17% worse on an unaided exam than peers who had never used it. This exposure is unlikely to be evenly spread – a risk judged especially acute for the novice school-age learners this guide addresses (Lodge & Loble, 2026). Students with weaker self-regulation are likely to lean on the tool in the way Fan et al. observed, deepening the gap rather than closing it, while stronger self-regulators, whose habits of monitoring and questioning their own understanding are well established (Zimmerman, 2000; Pintrich, 2004), are better placed to treat an AI response as a starting point to interrogate rather than an answer to accept.

This is not the whole picture, and the case should not be overstated beyond what the evidence yet shows. The same responsiveness that can erode effort could, well designed, sustain it – offering a student struggling with a concept a way back into a task, adapting to their level, or giving feedback when no one else is available to give it. In the same experiment, a version of the chatbot constrained to offer hints rather than answers removed the harm to learning entirely, although it produced no measurable gain over students who worked unaided (Bastani et al., 2025) – evidence that the outcome is shaped by how the tool is designed and used, not by the tool itself. The direct evidence on GenAI and motivation is emerging, and firm claims are premature. So the useful question is not whether GenAI helps or harms motivation in the abstract, but whether a particular task is designed so that the support it offers protects effort rather than removes it. The risk this section describes is real, and for some students likely – but it is one to design against, not a fixed outcome.

This connection between self-regulation and cognitive load is not incidental, and a growing body of research now brings the two together (de Bruin & van Merriënboer, 2017; Seufert, 2018). The central insight is that self-regulating is not free: monitoring your understanding, deciding what to do next, and adjusting your approach all draw on the same limited working memory that the learning task itself is using. This has two implications for GenAI:

First, when a task or a long, dense AI response overloads working memory, students have little capacity left to notice that they have stopped understanding, let alone act on it; the conditions for self-regulation break down at exactly the point they are most needed.

Second, because regulating your own learning is itself effortful, asking students to monitor and interrogate their GenAI use adds to the load, rather than removing it. That effort is worth protecting, but it has to be deliberately supported – especially for novices and students with weaker prior knowledge, who have the least spare capacity to spend on it. This is why Key Area 4 cannot be treated as an add-on to the other three. Metacognition and self-regulation are not additional skills students pick up after they have learned content, rather, they are the conditions under which the learning in Key Areas 1, 2, and 3 becomes effective. A student's job is not to manage cognitive load alone; that is first the teacher's responsibility. Without self-regulation, the benefits of teacher support, retrieval, and transfer are diminished. Self-regulation is not the final stage of learning, or something that comes only after content has been learned. It is engaged throughout – the thread that holds the other stages together (Seufert, 2018).



HOW GENAI CAN SUPPORT METACOGNITION AND SELF-REGULATION

Used thoughtfully, GenAI can potentially build these capabilities rather than erode them. The key is to design tasks that support students to monitor and evaluate their own thinking, rather than simply receive and reproduce the chatbot's.

GenAI may help students find the edge of their own understanding – the point where an explanation becomes uncertain, reasoning goes in circles, or confidence outruns their actual knowledge. When a student explains what they think they know and asks the chatbot to identify the gaps, they are doing exactly the self-monitoring that self-regulated learners do by habit.

GenAI can also support metacognitive planning – helping students think through how they will approach a complex task before starting: what they already know, where they expect to struggle, and what to do first. That kind of deliberate preparation is a hallmark of self-regulated learning, and one many students do not do on their own.

One conclusion has held across all four key areas of this guide: the educational value of GenAI is not fixed in the tool but determined by how teachers and students use it. Its value, or its cost, depends entirely on whether it is used in ways that preserve and strengthen the cognitive and metacognitive processes through which learning actually happens. Attention, knowledge building, transfer, and self-regulation are not a sequence that ends when a lesson does. They are the ongoing processes by which students become more capable, more independent learners. The teacher's role is not simply to permit or restrict access to AI. It is to design the conditions under which students are pushed to think, to notice their own thinking, and to take responsibility for both.

The following classroom example shows what this looks like when a teacher designs deliberately for metacognitive development, using GenAI as a diagnostic tool rather than a shortcut.

A note on context

The example uses World War I and the Australian home front as the curriculum content – a topic that appears in History courses across Australia. The pedagogical approach applied by Ms Ferreira on the following pages – the purpose-built agent, the constrained system prompt, the independent attempt before GenAI engagement – applies equally to other topics that require students to weigh the reliability and purpose of historical sources. The curriculum references in this example – the Australian Curriculum v9 content descriptions, historical skills, and Year 9 achievement standard – apply nationally; teachers can adapt the same approach to their own state or territory's curriculum or, with older students, to senior secondary frameworks.

CLASSROOM EXAMPLE: YEAR 9 HISTORY

NOTICING YOUR OWN THINKING

THE AGENT

Ms Ferreira teaches Year 9 History. Her class is partway through their depth study on World War I, focusing on the Australian home front. Over several weeks, she has built a purpose-designed AI agent – not a general-purpose chatbot, but a tool populated with specific curriculum knowledge:

What Ms Ferreira built into the agent

- Australian Curriculum v9 Year 9 History content descriptions for World War I (1914–1918) (Australian Curriculum, Assessment and Reporting Authority [ACARA], n.d.)
- The Year 9 historical skills, especially analysing the origin, content, context and purpose of primary and secondary sources
- The Year 9 History achievement standard
- The school's source-analysis rubric (origin, purpose, reliability and usefulness) with annotated sample responses at different levels
- A curated set of WWI sources – Australian recruitment posters, atrocity propaganda, and contextual scholarly material on wartime propaganda and the home front
- A system prompt instructing the agent never to assign a grade, never to write or rewrite student responses, to respond only with questions and criteria-referenced observations – and never to resolve a genuine historical debate for the student or present a contested interpretation as settled fact

The final instruction in the system prompt is the most important. Ms Ferreira has written it deliberately: "Your role is to help students identify what is missing or underdeveloped in their own work, using the source-analysis criteria as your reference. You do not assess. You do not rewrite. You ask. Where historians genuinely disagree, you show the student the debate; you do not resolve it for them."

SCIENCE OF LEARNING DESIGNING FOR SELF-REGULATION



The system prompt constraint – you do not assess, you do not rewrite, you ask – is not a technical limitation. It is a pedagogical decision. Ms Ferreira has designed the agent to refuse the shortcuts that would bypass student thinking. Every time the agent asks a question rather than providing an answer, it creates a moment where the student must evaluate the source’s purpose and reliability themselves, decide what it can and cannot support, and choose what to do next.

The quality of the knowledge the agent draws on – the marking rubric, the achievement standard, annotated exemplars, and scholarly context – means its questions are grounded in the actual standard and in real history, not a general approximation. This is what distinguishes a purposefully designed agent from a general chatbot: the constraints are as important as the content.

DIGITAL LITERACY CONTENT RISK AND CONSTRAINED DESIGN



Early and influential research on large language models showed that tools trained on web-scale data can absorb dominant patterns and framings present in their training sources (Bender et al., 2021). For a topic like WWI, a general chatbot may therefore reproduce simplified or incomplete historical framings rather than careful scholarly nuance. Recent Australian research found that chatbots can also present contested claims alongside legitimate information without distinguishing between them – a pattern the researchers describe as “bothsidesing rhetoric” (FitzGerald et al., 2026). Ms Ferreira’s purpose-built agent responds to that risk by refusing to resolve debates and by grounding its questions in scholarly sources and the curriculum content, pushing students to distinguish what is contested from what is settled.

WWI is less commonly associated with overt conspiracy and distortion than some other twentieth-century subjects, so the rationale here is more about contestability and source provenance rather than content safety. Even so, war content is emotive, and some students will have family service connections. The agent does not replace the teacher’s judgement about how to handle that human dimension with care.

BEFORE THE AGENT - INDEPENDENT ATTEMPT FIRST

Ms Ferreira does not begin the lesson with the agent. She displays a source on the board – an Australian recruitment poster from 1915, showing an idealised soldier and a patriotic call to enlist – and gives the class ten minutes of silence.

Source Analysis Task

Study Source A, an Australian recruitment poster from 1915. Explain what this source reveals about how the Australian government tried to persuade men to enlist.

MS FERREIRA

"Write your response on your own. No notes and no agent. You have ten minutes. When you are done, annotate your own response – mark the places where you feel confident and put a question mark next to anything you are unsure of."

Students write. The self-annotation step is deliberate – Ms Ferreira uses it to make students locate their own uncertainty before the agent does it for them. After ten minutes, students open the agent and submit their response. Ms Ferreira instructed the students to add to the end of their response: Check this answer against the criteria.

SCIENCE OF LEARNING METACOGNITIVE PREPARATION - LOCATING UNCERTAINTY BEFORE FEEDBACK ARRIVES



The self-annotation step is a deliberate metacognitive act. Before any external feedback arrives, students are asked to identify where their confidence is and where it is not warranted – to monitor their own understanding and mark it on the page. This matters because students who have already located their own uncertainty engage with feedback differently from students who have not. Rather than receiving the agent's questions as a general challenge, they can connect them to the specific places they already suspected were weak.

The independent attempt also preserves the retrieval effort that builds durable knowledge. Students must produce what they can from their own understanding of the source before they use the agent to refine it – the thinking happens first, and then the AI refines it.

STAGE ONE - SOURCE ANALYSIS

Two students' exchanges with the agent illustrate how it responds differently depending on what is brought to it.

DANIEL'S EXCHANGE WITH THE AGENT

Daniel

Here is my response: "This source shows that lots of Australians wanted to join the army, because the poster makes war look heroic and exciting. It uses a strong-looking soldier and bright colours to make people feel patriotic. This proves Australians were keen to enlist." Please check it against the criteria.

AI chatbot

You've spotted the techniques clearly – the heroic figure, the patriotic appeal, the use of colour. That's a good start, and the criteria reward identifying how a source tries to persuade. A few questions worth sitting with. This is a recruitment poster: it was made by the government to persuade men to enlist. Does it tell us that Australians were keen, or that the government wanted them to be? Those are different claims. And here's one to think about: if everyone was already eager to sign up, why would the government spend money producing posters to persuade them?

Daniel

So it doesn't actually prove people wanted to join. It shows the government was trying to make them want to join. Maybe they made it because not enough people were signing up?

AI chatbot

That's a great inference. The existence of heavy recruitment propaganda can suggest that enlistment was a concern – it does not prove that enthusiasm was universal. To find out whether Australians were actually keen, what other evidence would you look for? Think about what kind of source would record how people really responded, rather than how the government wanted them to respond. Then have a go at revising your last sentence so it says what the source can actually support.

Daniel

I could look at enlistment numbers over time, or letters and diaries, or the conscription votes. My new sentence: "Source A reveals the persuasive techniques the government used to encourage enlistment like heroic imagery and patriotic appeals, which suggests the government was working hard to boost recruitment, rather than proving Australians were already eager to enlist."

AI chatbot

That sentence is significantly stronger. It matches what the source can actually support. One last thing: your first version used the word 'proves'. Historians are careful with that word. Does a single recruitment poster 'prove' anything about public opinion, or does it 'suggest' and 'reveal'? Notice how changing that one word changed how defensible your sentence is.



SCIENCE OF LEARNING SELF-MONITORING THROUGH ELABORATIVE DIALOGUE

Daniel's first response treated the poster as a window into what Australians thought – which is quite common and understandable. The agent did not correct him; it interrogated the claim. Each question required Daniel to do something specific: separate what the source is trying to do from what it proves, infer why the source exists, and then evaluate the precision of his own language.

The shift Daniel makes – from 'this shows what people thought' to 'this shows what the government wanted people to think' – is the foundation of source analysis. He is learning to read a source as evidence of its own making, not as a neutral record of events. That is a more durable understanding than a list of techniques, and he constructed it himself.

ISABELLE'S EXCHANGE WITH THE AGENT

Isabelle

My response: "Source A is a government recruitment poster, so it is designed to persuade rather than to record. It reveals the values the government appealed to masculinity, patriotism, duty and mateship and shows how the war was sold to the public. But because it is propaganda, it shows the message the government wanted to project, not how the public actually responded. To understand attitudes to the war I would need to compare it with enlistment data and the conscription referendum results." Can you check this against the criteria?

AI chatbot

This is a strong response – you've identified the source's purpose and its limitation, which is exactly what the higher levels of the rubric reward. Let me push you further. A source can be unreliable as evidence of one thing while being excellent evidence of another. The poster is weak evidence of public opinion, but what is it strong evidence of? And here is a tension worth using: the 1916 and 1917 conscription referendums were both defeated. What does that gap – between the poster's confident patriotic message and a public that twice voted against compulsion – suggest about wartime Australia?

Isabelle

So the poster is reliable evidence of the government's intended message, and of the values it assumed would motivate men – that's historically revealing in itself. And the referendums failing shows the public was divided, even though the posters present enlistment as the obvious patriotic choice. The gap between the message and the vote is actually the interesting part as it shows the values in the poster didn't persuade everyone.

AI chatbot

That's a sophisticated move. You've turned the source's limitations into evidence. Using a source to understand the assumptions and anxieties of the people who made it is a higher-order skill than using it as a record of events. How would you express that distinction – reliable for the government's message, limited for public opinion – in a single sentence, without overclaiming in either direction?



SCIENCE OF LEARNING SELF-REGULATION AT THE EDGE OF COMPETENCE

Isabelle's response was already strong: she identified purpose and limitation without prompting. The agent's role here is not to fill gaps but to push her to notice something she had not yet considered. This is a higher-order metacognitive move: pushing her to see that a source's unreliability for one purpose can make it valuable for another, and to use the tension between the poster and the referendum results rather than ignore it.

The insight she reaches – that the gap between the official message and the public vote is the interesting part – is genuinely new understanding, not a retrieval of something she already knew. The agent's grounding in the rubric and the actual history is what makes this challenge precise rather than generic: it can ask the question that moves a capable student from a good answer to a more analytical one.



DIGITAL LITERACY SOURCE CREDIBILITY AND HISTORICAL EVIDENCE

The habit Daniel and Isabelle are building – asking not just what the source says but what it is for and therefore can reliably reveal – is the same critical habit that productive GenAI use requires. A propaganda poster and an AI response share an important feature: fluency, confidence, and wide circulation are not evidence of accuracy. WWI contains material that was widely believed at the time and later shown to be fabricated or exaggerated, including atrocity stories and faked photographs. Students who learn to interrogate a source's purpose in history are building the same transferable critical skill they need to evaluate a confident, fluent AI answer.

MS FERREIRA PAUSES THE CLASS

After fifteen minutes, Ms Ferreira stops the class before anyone moves to the extended response.

MS FERREIRA

"Close the agent. Look at your original response and what you have now. Write two sentences: what did the agent ask you that you hadn't asked yourself? And what does that question tell you about what you need to keep working on?"

The room shifts to writing. Ms Ferreira moves between desks reading responses. She is not commenting on individual work – she is looking for patterns. She notices that several students, like Daniel, had treated the poster as a transparent record of what people thought, rather than as evidence of what the government wanted them to think. She makes a note to address the difference between a source's purpose and its content directly in next week's lesson.

SCIENCE OF LEARNING METACOGNITIVE AWARENESS AND TEACHER INSIGHT



The two-sentence reflection serves two purposes at once. For students, it requires them to articulate what the agent revealed about their thinking – a metacognitive step that makes the gap between current understanding and the expected standard visible and explicit. Naming that gap is the first step toward closing it.

For Ms Ferreira, patterns in these reflections provide insight that marking paragraphs alone could not: not just what students got wrong, but what conceptual distinction they are not yet making. The agent has effectively run a class set of individual diagnostic conversations in parallel – and Ms Ferreira can read the results in a few minutes. This is the interaction log as formative assessment: each student's exchange is a visible record of their thinking, giving the teacher information without substituting for her own diagnostic eye.

STAGE TWO - EXTENDED RESPONSE

Ms Ferreira now moves the class to a short extended response. She is explicit about how the agent's role changes.

MS FERREIRA

"For the extended response, the agent will not check your finished paragraph. It will help you stress-test your argument before you write. You are going to use it to interrogate your main claim and your evidence plan – not your finished product. Write the paragraph yourself."

Extended Response Question

How useful is Source A as evidence of Australian attitudes to the war? In your answer, refer to the source's origin, purpose and content, and to at least one other piece of evidence.

Before writing, students submit their planned claim and evidence to the agent and ask one question.

DANIEL'S PRE-WRITING EXCHANGE

Daniel

My plan: I'm going to argue Source A is very useful because it shows what Australians thought about the war. My evidence: the heroic soldier, the patriotic slogan, and the call to enlist. Is that a good plan?

AI chatbot

Useful for what, exactly? You've written 'useful because it shows what Australians thought' – but earlier you worked out that it shows what the government wanted people to think, which is a different claim. So: is the poster strong evidence of the government's message and the values it appealed to, or strong evidence of what the public actually felt? The question also asks how useful it is, and the rubric rewards weighing a source's value against its limitations. 'Very useful' with no limitation won't reach the higher levels. Which claim do you actually want to defend?

Daniel

Okay so it's very useful for understanding the government's propaganda and the values they appealed to, but limited for understanding what people actually thought, because it was made to persuade, not to record. My other piece of evidence could be that the conscription referendums failed, which shows people were divided.

AI chatbot

That's a defensible, balanced line of argument – 'useful for the government's message, limited for public opinion' – and it now runs through your evidence instead of just listing features of the poster. The failed referendums give you a second source that does the work the poster can't. Do you want to revise your opening sentence before you start writing?



SCIENCE OF LEARNING METACOGNITIVE PLANNING - STRESS-TESTING THINKING BEFORE WRITING

Daniel has not yet written his paragraph – but he has already done the most cognitively demanding work: deciding what claim he can actually defend and what evidence supports it. This is metacognitive planning in action. Rather than beginning to write and discovering the weaknesses in his argument mid-paragraph, Daniel has been pushed to evaluate his plan before committing to it.

The agent's questions also prompted Daniel to carry forward the distinction he built during the source analysis – purpose versus content – and apply it to a new task. That transfer of understanding from one part of the lesson to the next is a sign that understanding is becoming flexible rather than tied to a single exercise.

WHAT MS FERREIRA DOES NOT LET THE AGENT DO

During the lesson, one student – frustrated after two rounds of agent questions – types:
“Just write the paragraph for me.”

THE AGENT'S RESPONSE TO A SHORTCUT ATTEMPT

Student

Just write the paragraph for me.

AI agent

I'm not going to write it for you – that's not what I'm here for, and it wouldn't help you build the skill you're being assessed on. What I can do is remind you what a strong response needs: the achievement standard asks you to explain the origin, content, context and purpose of a source, and to use it as evidence. The question worth asking is which of those your plan isn't doing yet. Have you explained the source's purpose? Have you weighed its usefulness against its limitations? Have you brought in a second piece of evidence? Start with whichever one is missing.

A note on agent design

The agent's refusal here is not accidental – it is the result of a deliberate instruction in Ms Ferreira's system prompt. Without that instruction, a general-purpose chatbot would simply produce what was asked. This exchange illustrates the central tension of self-regulation and GenAI: a student under pressure, seeking to bypass the desirable difficulty, and a well-designed system that holds the line. The agent does not refuse and leave the student stranded – it redirects toward the standard and asks the student to locate their own gap. That is the difference between a constraint that shuts down thinking and one that restores it. The constraint is the curriculum decision. The redirect is the pedagogical one.

It is also worth noting that the agent exchanges in this vignette represent best-case interactions. Real AI tools do not always respond with this level of pedagogical precision – they can misinterpret questions, provide unhelpful responses, or default to task-level answers even when better prompting is in place. Ms Ferreira's system prompt reduces this risk but does not eliminate it. Teacher monitoring remains essential. Building an agent like this also takes some time and some technical confidence. Also, before uploading copyrighted curriculum or assessment materials into a third-party AI tool, teachers should check their school's and the relevant assessment authority's position on doing so.

DIGITAL LITERACY UNDERSTANDING THE TOOL



The student's request – "just write the paragraph for me" – reflects a misunderstanding not only of learning but also of the agent's purpose. Using any digital tool well requires understanding both its purpose and its limitations. Ms Ferreira's system prompt makes those limitations explicit by design. But in classrooms where tools are not purpose-built, students need to develop that critical awareness themselves by asking not just what the tool can do, but what it should be asked to do, and why. This may include re-prompting, asking for more sources, or checking in with the teacher.

A FINAL WORD

The four key areas explored in this guide – attention and cognitive load, memory and knowledge, transfer and problem-solving, and metacognition and self-regulation – are not a sequence that ends when a lesson finishes.

They are the ongoing processes through which students become increasingly capable, independent learners. Each depends on the others. Attention without knowledge goes nowhere. Knowledge without transfer stays where it is or is forgotten. Transfer without self-regulation is fragile. And none of it is sustained without the motivation and awareness to keep going when learning is hard.

GenAI does not change any of this. What it changes is the ease with which students can appear to be doing the cognitive work without actually doing it – and the ease with which teachers can mistake that appearance for the real thing. That is the educational challenge this guide has named clearly.

The teachers in these vignettes – Ms Orlando, Mr Chen, Ms Clifford, Ms Ferreira – are not technology experts. They are expert teachers who understand how learning works and have thought carefully about how a new tool fits into that understanding.

THEY ASK GOOD QUESTIONS. THEY WATCH THEIR STUDENTS CLOSELY. THEY INTERVENE AT THE LEVEL OF THINKING, NOT JUST BEHAVIOUR. THEY PROTECT THE MOMENTS OF DESIRABLE DIFFICULTY THAT BUILD REAL UNDERSTANDING. AND THEY USE GENAI WHEN IT GENUINELY HELPS AND HOLD IT BACK WHEN IT WOULD GET IN THE WAY.

That is not a technology question. It is a teaching question. And it is the right question to be asking.

The science of learning does not tell teachers exactly what to do with GenAI. It tells them what to pay attention to – the cognitive processes that make learning stick, the conditions that support them, and the risks that undermine them. This guide makes that knowledge practical and specific enough to be genuinely useful in classrooms.



REFERENCES

- Anderson, R. C., Reynolds, R. E., Schallert, D. L., & Goetz, E. T. (1977). Frameworks for comprehending discourse. *American Educational Research Journal*, 14(4), 367–381. <https://doi.org/10.3102/00028312014004367>
- Australian Curriculum, Assessment and Reporting Authority. (n.d.). *History 7–10: Curriculum content*. Australian Curriculum.
- Bandura, A. (1997). *Self-efficacy: The exercise of control*. W. H. Freeman.
- Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge University Press.
- 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), Article e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FACCT '21)*, 610–623. <https://doi.org/10.1145/3442188.3445922>
- Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In M. A. Gernsbacher, R. W. Pew, L. M. Hough, & J. R. Pomerantz (Eds.), *Psychology and the real world: Essays illustrating fundamental contributions to society* (pp. 56–64). Worth Publishers.
- Bjork, R. A., Dunlosky, J., & Kornell, N. (2013). Self-regulated learning: Beliefs, techniques, and illusions. *Annual Review of Psychology*, 64, 417–444. <https://doi.org/10.1146/annurev-psych-113011-143823>
- Blakemore, S.-J. (2018). *Inventing ourselves: The secret life of the teenage brain*. Doubleday.
- Blakemore, S. J., & Choudhury, S. (2006). Development of the adolescent brain: Implications for executive function and social cognition. *Journal of Child Psychology and Psychiatry*, 47(3–4), 296–312. <https://doi.org/10.1111/j.1469-7610.2006.01611.x>
- Brown, A. L. (1980). Metacognitive development and reading. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp. 453–481). Erlbaum.
- Cepeda, N. J., Pashler, H., Vul, E., Wixted, J. T., & Rohrer, D. (2006). Distributed practice in verbal recall tasks: A review and quantitative synthesis. *Psychological Bulletin*, 132(3), 354–380. <https://doi.org/10.1037/0033-2909.132.3.354>
- Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, 13(2), 145–182. https://doi.org/10.1207/s15516709cog1302_1
- Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5(2), 121–152. https://doi.org/10.1207/s15516709cog0502_2
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1), 87–114. <https://doi.org/10.1017/S0140525X01003922>
- Cowan, N. (2010). The magical mystery four: How is working memory capacity limited, and why? *Current Directions in Psychological Science*, 19(1), 51–57. <https://doi.org/10.1177/0963721409359277>
- Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, 11(6), 671–684. [https://doi.org/10.1016/S0022-5371\(72\)80001-X](https://doi.org/10.1016/S0022-5371(72)80001-X)
- de Bruin, A. B., & van Merriënboer, J. J. (2017). Bridging Cognitive Load and Self-Regulated Learning Research: A complementary approach to contemporary issues in educational research. *Learning and Instruction*, 51, 1–9. <https://doi.org/10.1016/j.learninstruc.2017.06.001>
- Dehaene, S. (2020). *How we learn: The new science of education and the brain*. Viking.
- 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>
- Feldon, D. F. (2007). Cognitive load and classroom teaching: The double-edged sword of automaticity. *Educational Psychologist*, 42(3), 123–137. <https://doi.org/10.1080/00461520701416173>
- FitzGerald, K. M., Riedlinger, M., Bruns, A., Harrington, S., Graham, T., & Angus, D. (2026). *Just asking questions: Doing our own research on conspiratorial ideation by generative AI chatbots*. *Media and Communication*, 14, Article 11337. <https://doi.org/10.17645/mac.11337>

- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, 34(10), 906–911. <https://doi.org/10.1037/0003-066X.34.10.906>
- Hattie, J., & Timperley, H. (2007). The power of feedback. *Review of Educational Research*, 77(1), 81–112. <https://doi.org/10.3102/O03465430298487>
- Hutchins, E. (1995). *Cognition in the wild*. MIT Press.
- 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
- Kornell, N., & Bjork, R. A. (2008). Learning concepts and categories: Is spacing the “enemy of induction”? *Psychological Science*, 19(6), 585–592. <https://doi.org/10.1111/j.1467-9280.2008.02127.x>
- Kuhn, D. (2000). Metacognitive development. *Current Directions in Psychological Science*, 9(5), 178–181. <https://doi.org/10.1111/1467-8721.00088>
- Lodge, J. M., & Loble, L. (2026). *Artificial intelligence, cognitive offloading and implications for education*. University of Technology Sydney. https://figshare.uts.edu.au/articles/report/Artificial_intelligence_cognitive_offloading_and_implications_for_education/31302475
- Panadero, E. (2017). A review of self-regulated learning: Six models and four directions for research. *Frontiers in Psychology*, 8, 422. <https://doi.org/10.3389/fpsyg.2017.00422>
- Pintrich, P. R. (2004). A conceptual framework for assessing motivation and self-regulated learning in college students. *Educational Psychology Review*, 16(4), 385–407. <https://doi.org/10.1007/s10648-004-0006-x>
- Renkl, A. (2014). Toward an instructionally oriented theory of example-based learning. *Cognitive Science*, 38(1), 1–37. <https://doi.org/10.1111/cogs.12086>
- Roediger, H. L., & 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>
- Rosenshine, B. (2012). Principles of instruction: Research-based strategies that all teachers should know. *American Educator*, 36(1), 12–19.
- Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55(1), 68–78. <https://doi.org/10.1037/0003-066X.55.1.68>
- Seufert, T. (2018). The interplay between self-regulation in learning and cognitive load. *Educational Research Review*, 24, 116–129. <https://doi.org/10.1016/j.edurev.2018.03.004>
- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4
- Sweller, J. (2006). The worked example effect and human cognition. *Learning and Instruction*, 16(2), 165–169. <https://doi.org/10.1016/j.learninstruc.2006.02.005>
- Sweller, J. (2011). Cognitive load theory. In J. P. Mestre & B. H. Ross (Eds.), *Psychology of learning and motivation* (Vol. 55, pp. 37–76). Academic Press. <https://doi.org/10.1016/B978-0-12-387691-1.00002-8>
- Sweller, J., van Merriënboer, J. J. G., & Paas, F. (2019). Cognitive architecture and instructional design: 20 years later. *Educational Psychology Review*, 31(2), 261–292. <https://doi.org/10.1007/s10648-019-09465-5>
- van Merriënboer, J. J. G., & Sweller, J. (2005). Cognitive load theory and complex learning: Recent developments and future directions. *Educational Psychology Review*, 17(2), 147–177. <https://doi.org/10.1007/s10648-005-3951-0>
- 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.
- Zimmerman, B. J. (2000). Attaining self-regulation: A social cognitive perspective. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds.), *Handbook of self-regulation* (pp. 13–39). Academic Press. <https://doi.org/10.1016/B978-012109890-2/50031-7>

