What do we still know when the tool is switched off?

Deep learning has an older meaning than the neural network. It is how people build the understanding their judgement runs on, and it is the step AI makes easiest to skip.

Sylvain Rottier
Co-Founder, Mindmarqs | Psychologist
October 2026

In “What is the point of healthy brains when we do not use them?” I called the third behaviour of cognition management active learning. I now call it deep learning, after the research it comes from, even though in a room of managers half the audience will hear neural networks. The name is worth the confusion. Long before computers learned anything, deep learning described how people do.

In 1976 Ference Marton and Roger Säljö asked students to read an academic text and then questioned them about it. Some had read to understand what the author meant. Others had read to be able to reproduce it. The researchers called these deep and surface processing, and the difference showed in what the students had understood afterwards. It lay in what they were trying to do while they read.

That distinction matters more now than it did then, because we have built a machine that is very good at surface.

Fluent is not the same as understood

An AI summary is fluent, organised and confident. Reading it feels like understanding. It is not the same thing, and the gap only shows when the tool is gone.

Some of the clearest evidence comes from experts using AI in their real work. In four endoscopy centres in Poland, nineteen experienced doctors, each with more than 2,000 colonoscopies behind them, began using AI that highlights possible growths during the procedure. In the months that followed, in procedures done without the AI, their detection rate fell from 28.4% to 22.4% (Budzyń and colleagues, 2025). It is an observational study, and one study. But these were experts, and the skill that slipped was the one they are trusted for.

A randomised experiment points at the mechanism. Fifty-two software engineers learned a programming library they had never used, half with an AI assistant and half without. On a test of what they had understood, the AI group scored 50% against 67% (Shen and Tamkin, 2026, a preprint). The more interesting finding was inside the AI group. Those who handed the work over, or leaned on the assistant more and more, scored lowest. Those who asked it to explain the code it wrote, or used it only for conceptual questions and solved the problems themselves, scored much better.

Same tool, two outcomes. The difference was whether the person was still doing some of the thinking.

Two kinds of offloading

Handing thinking to tools is not new, and it is not wrong. We write things down, we use calculators, we search instead of memorising. Evan Risko and Sam Gilbert, reviewing the research on what they call cognitive offloading, describe it as often useful and rational, and they are clear about what it does: it moves what is remembered from the person to the device.

So the question is not whether to offload, but what. Offloading the parts we do not need to own, the formatting, the first search, the reference we can look up again, frees attention for the parts we do. Offloading the parts that build understanding, working out why, connecting it to what we already know, explaining it, leaves us with an answer and nothing underneath it. The first is judgement. The second is surrender, and it rarely feels like a decision at the time.

Why this matters for decisions

Judgement runs on context: what happened last time, what the numbers usually look like, which assumptions held and which did not. That context is not downloaded. It is built, a little at a time, every time we do some of the mental work ourselves.

Organisations depend on it more than they notice. A decision that nobody really worked through is hard to stand behind and harder to explain six months later. When the people in the room cannot say why a choice was made, the organisation has the decision but not the understanding, and the next decision starts from zero.

What builds understanding

The learning science here is mature, and it agrees on a simple principle: we understand what we construct, not what we receive.

Michelene Chi and Ruth Wylie brought decades of studies together in what they call the ICAP framework. Learning improves as engagement moves from passive, receiving the material, to active, doing something with it such as highlighting, to constructive, producing something new such as an explanation or a diagram, to interactive, building on each other's thinking. Reading an AI summary is passive. Rewriting it in your own words is constructive.

Explaining is one of the strongest tools we have. A meta-analysis of 64 studies with nearly 6,000 participants found that prompting people to explain material to themselves had an average effect of $g = 0.55$, a substantial gain, and that it did better than being given the explanation by an instructor (Bisra and colleagues, 2018). The answer we generate teaches us more than the answer we are handed.

Drawing works too. In experiments at the University of Waterloo, people who drew the objects on a word list, even badly and in a few seconds, often recalled more than twice as many words as people who wrote them out (Wammes, Meade and Fernandes, 2016). A sketch forces you to decide what the thing actually is.

And the habits most of us rely on are the weakest ones. In a review of ten common learning techniques, Dunlosky and colleagues rated practice testing and spaced practice as the most useful, and rereading and highlighting among the least. At work we reread the deck and highlight the report. We rarely test ourselves on what it said.

What deep learning looks like at work

None of this means going back to school. It means doing a small part of the thinking yourself, at the moments that build understanding.

1. Think first, then ask. Write two or three lines of your own view before you open the chat window. You will read the answer differently, and you will notice when it disagrees with you.
2. Ask for the why, not only the what. When AI gives you an answer, ask it to explain its reasoning, and then check whether you could explain it yourself.
3. Explain it back. Before you sign off a recommendation, say in your own words why it is right. If you cannot, you have found the gap.
4. Sketch it. One rough drawing of how the thing works, on paper or a whiteboard, does more than another read of the slides.
5. Test, do not reread. After an important meeting, write down the three things that mattered before you look at the notes.
6. Come back to it. A week after a significant decision, take five minutes to ask what you have learned since. Spacing turns an event into experience.

Each takes minutes. Together they are the difference between using a tool and being carried by it.

These are habits, and habits can be trained

Deep learning is not a talent. The techniques above work because they change what people do, and they can be prompted, practised and built into the working week. In the self-explanation studies, people were not told what to think, only asked to explain. The benefit came from the act itself.

This is how we approach it at Mindmarqs. Building context and memory is one of the three behaviours our trainings work on. We start from each person's own pattern, practise on the decisions they actually face, and tie the new habits to moments in their week so that they hold: see it, steer it, keep it.

One decision this week

Take one decision you will face this week, and ask three questions:

- Before: what do I already think, before I ask anyone or anything?
- During: could I explain this recommendation in my own words?
- After: what do I understand now that I did not before?

Use the tool. Use it a lot. Just make sure that when it is switched off, the understanding is still yours.

Sources

- Bisra, K., Liu, Q., Nesbit, J. C., Salimi, F. & Winne, P. H. (2018). Inducing self-explanation: A meta-analysis. Educational Psychology Review, 30(3), 703–725.
- Budzyń, K., Romańczyk, M., Kitala, D. et al. (2025). Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study. The Lancet Gastroenterology & Hepatology, 10(10), 896–903.
- Chi, M. T. H. & Wylie, R. (2014). The ICAP framework: Linking cognitive engagement to active learning outcomes. Educational Psychologist, 49(4), 219–243.
- Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J. & Willingham, D. T. (2013). Improving students' learning with effective learning techniques. Psychological Science in the Public Interest, 14(1), 4–58.

- Marton, F. & Säljö, R. (1976). On qualitative differences in learning: I. Outcome and process. British Journal of Educational Psychology, 46(1), 4–11.
- Risko, E. F. & Gilbert, S. J. (2016). Cognitive offloading. Trends in Cognitive Sciences, 20(9), 676–688.
- Shen, J. H. & Tamkin, A. (2026). How AI impacts skill formation. arXiv preprint 2601.20245.
- Wammes, J. D., Meade, M. E. & Fernandes, M. A. (2016). The drawing effect: Evidence for reliable and robust memory benefits in free recall. Quarterly Journal of Experimental Psychology, 69(9), 1752–1776.