Outsourcing Our Minds: AI, Cognitive Offloading, and the Future of Thinking

Dr Neil Hopkin
Director of Education
Fortes Education

The Paradox of Intelligence: When Machines Think, Do We Stop?

On the morning of 5 January 2011, Kevin Poulsen was sitting at his desk, staring at his screen. He was reading something—an article about memory, or at least that's what he thought it was about. The piece, written by Betsy Sparrow and her colleagues at Columbia University, contained a deceptively simple experiment: participants were given trivia statements to read, some easy, some difficult, and then later asked to recall them. It was the sort of thing that cognitive psychologists had been doing for decades, but with one crucial difference. Half the participants were told the statements would be erased from the computer, while the others were assured the information would be saved.

You can probably guess what happened. Those who thought they'd have no external record tried harder, engaged more, and remembered more. Those who believed they could look it up later didn't even bother. The internet, it seemed, had done something strange to memory—not just changing what we remember, but how we decide what's worth remembering at all (Sparrow, Liu & Wegner, 2011).

Now, extend that logic beyond trivia facts. What happens when this same principle applies not just to names and dates, but to reasoning, decision-making, and problem-solving? And what will happen when the machine that remembers also begins to think?

The Rise of External Intelligence

There's a tendency to believe that new technologies free us. The printing press liberated knowledge from monasteries. The pocket calculator freed schoolchildren from the tyranny of long division. Google search eliminated the need for encyclopaedias, and ChatGPT can now explain quantum physics in the voice of Oscar Wilde. At every stage, tools have promised efficiency, promising to do the hard work for us so that we can do something else, something 'more valuable'.

But as Nicholas Carr (2010) pointed out in *The Shallows*, that bargain has always come with a cost. Each technological advance changes not just what we do, but how we think. When the ancient Greeks introduced writing, Socrates worried that it would ruin memory, turning knowledge into something people referenced rather than truly understood. Gutenberg's printing press was accused of making scholars lazy, because why memorise Aristotle when you could now *own* him? The concern was never about information itself, but about what happened to the mind when it no longer had to work to retrieve it.

And then came AI.

AI doesn't just store knowledge; it generates it. It doesn't just provide answers; it anticipates questions. And therein lies the paradox: the smarter our tools become, the less we have to think. The less we have to think, the more we rely on the tools. And the more we rely on the tools, the less we remember how to think at all. The term for this phenomenon is **cognitive offloading**, and it's not new. Risko and Gilbert (2016) describe it as the process by which people delegate cognitive tasks to the environment, reducing the mental effort required for problem-solving. Your phone remembers phone numbers so you don't have to. Google suggests the next word in your sentence before you've even finished typing. At first glance, this seems beneficial—why waste brainpower on trivial tasks when an external tool can handle them?

But here's the problem: cognition doesn't work like a ledger, with fixed deposits and withdrawals. Memory, reasoning, and creativity aren't separate processes that can be outsourced one by one without consequence. They are interconnected, and the more we offload, the weaker those connections become. Henkel (2014) found something fascinating when studying tourists at a museum. When visitors took photographs of an object, they were later less able to remember details about it. The act of taking a picture, of outsourcing the task of remembering, had altered the very way they engaged with the experience. The brain had decided it didn't need to store that information because the camera was doing the job instead. Go to any concert or sporting event nowadays, or even watch a parent observing their child's first steps or first birthday: the experience is often lived through the screen of the phone videoing it, rather than engaging in just 'being there'.

Now imagine the same effect on mathematics, science, and problem-solving. The rise of AI tutors, instant computational tools, and generative reasoning models means students no longer need to work through problems; they can ask ChatGPT to do it for them. This is not just about laziness. It's about *how* knowledge is acquired in the first place. Mills and Keil (2004) describe something called the **Illusion of Explanatory Depth:** the tendency to believe we understand complex concepts better than we actually do. The more readily available external explanations become, the stronger this illusion grows. You've probably experienced this yourself. You're in a conversation about politics, or astrophysics, the history of Impressionist art, or who the leading man was in *that* film. Someone makes a claim, and instead of debating, someone pulls out their phone. Within seconds, the discussion is over. The answer is retrieved, the ambiguity eliminated. But is that the same as *understanding*?

Ward et al. (2017) demonstrated that the mere presence of a smartphone—turned off, face-down, completely inert—reduces cognitive capacity. The brain, aware that information is just a search away, disengages from deeper processing. Fisher, Goddu and Keil (2015) found something even more striking: when people use the internet to answer a question, they later overestimate how much they know even about *related* topics. Access to AI-enhanced knowledge creates the *illusion* of intelligence rather than actual expertise. And this is where AI changes the game entirely. When ChatGPT produces an answer, it doesn't just provide information; it constructs a persuasive, articulate response that *feels* authoritative. The problem isn't that AI gets things wrong (although it does). The problem is that people trust it even when it does. Zaphir et al. (2024) argue that AI lacks genuine critical thinking, yet because it mimics the patterns of human reasoning, we assume it possesses something akin to wisdom.

The philosopher Evgeny Morozov (2013) calls this **technological solutionism**—the belief that for every problem, there is a software fix. Can't write an essay? AI will draft it. Struggling with an equation? AI will solve it. But what gets lost in the process? What do we trade away, invisibly, in each of these micro-exchanges of convenience? There was a time when chess grandmasters trained by poring over books of past games, reconstructing moves, absorbing strategy through mental effort. Today, they train with AI opponents who predict and suggest optimal moves instantly. The result? Players are stronger than ever, but many struggle to explain *why* certain strategies work. The thinking process, once slow and deliberate, has been altered. The machine does the heavy lifting; the human plays catch-up.

Science,mathematics, engineering—these fields depend not just on answers, but on *how* those answers are reached. When AI short-circuits that process, we risk becoming brilliant without *understanding*, efficient but not *insightful*. There is no undoing AI. The genie, as they say, is out of the bottle. But if we are to understand where we go from here, we must first understand what's at stake. Intelligence, after all, is not about having information but about *how* we engage with it.

The Rise of AI Tutors: Personalisation or Dependency?

In 1984, the educational psychologist Benjamin Bloom published what would become one of the most cited papers in the history of learning science. He had uncovered something that, on the surface, seemed almost too good to be true: students who received one-on-one tutoring outperformed those in conventional classroom settings by a staggering margin. The average tutored student scored two standard deviations higher than their peers—meaning that with a personal tutor, the average student could perform better than 98% of those in a traditional classroom (Bloom, 1984). The finding became known as the **2 Sigma Problem**, and it presented an extraordinary challenge. If tutoring was this effective, then the logical conclusion was simple: every student should have a personal tutor. But in practice, this was impossible. One-on-one tutoring was expensive, logistically unfeasible, and beyond the reach of most education systems. And so, the problem remained unsolved. Until now.

In the last decade, AI tutors have emerged as the first real attempt to crack the 2 Sigma Problem. They are adaptive, infinitely scalable, and capable of providing instant feedback. Unlike human tutors, they do not get tired, impatient, or distracted. A struggling student in a remote village has, in theory, access to the same level of personalisation as a high-achieving student in an elite private school. For the first time, the utopian vision of individualised learning for all feels within reach. But something is missing from this triumphant narrative.

In a study at the University of Oxford, researchers observed that students who used AI tutors for maths problem-solving initially showed higher success rates than those who relied on traditional methods. The AI provided instant correction, step-by-step guidance, and explanations tailored to each student's needs (Luckin et al., 2021). But when those same students were later asked to solve problems without AI assistance, their performance plummeted. They had learned to follow patterns, not to think through problems. The tutor had made them better at solving familiar problems but had done little to prepare them for situations where the solutions were not so easily accessible. The paradox of AI tutoring is not that it fails to teach. It teaches too well—so well that it risks making the learner redundant in the process.

When OpenAI released ChatGPT-4, one of the most immediate applications was education. Within months, millions of students around the world were using it to answer homework questions, draft essays, and clarify concepts they didn't understand. The AI could adapt its responses based on how much background knowledge a student had, explaining Newton's laws in simple terms to a ten-year-old or providing a university-level derivation of Maxwell's equations to an engineering student. This kind of responsiveness made AI feel not like a tool, but like a teacher. It mimicked the experience of learning from a human tutor—except with perfect patience and infinite availability. But something strange started happening.

Researchers in Singapore conducted an experiment where students were divided into two groups. One group used ChatGPT for real-time feedback while writing an essay. The other group relied on traditional peer review and teacher feedback. The students who used AI assistance wrote more polished, structurally coherent essays (Kovanović et al., 2023). Their grammar was cleaner, their arguments smoother. But when asked to revise and improve their writing independently, they struggled significantly more than their peers. The AI had acted as an invisible cognitive prosthetic—helping them without them realising it was helping them.

This aligns with what cognitive scientists have long suspected: when feedback is too immediate, too precise, it creates a form of passive learning (Kraut & Resnick, 2023). The effortful process of struggling through a concept, making mistakes, and finding one's own path to understanding is where real learning happens. AI short-circuits that process. It makes learning frictionless—which, paradoxically, may be the biggest problem of all.

The phenomenon has historical precedent. In the 1970s, educational psychologists discovered that children who used calculators too early in their schooling often failed to develop strong numerical reasoning skills (Fuson & Briars, 1979). They could compute quickly, but when asked to explain the relationships between numbers, they were lost. The device had enabled efficiency but at the cost of conceptual depth. A similar concern emerged in the 1990s with spellcheckers. Students who relied on autocorrect tools developed weaker spelling abilities, because they never internalised the rules they were bypassing (Anson & Schwegler, 1997).

Now, AI tutors are extending this cognitive outsourcing to reasoning itself. A study from Stanford University found that students using AI-assisted maths solvers developed a false sense of confidence in their problem-solving abilities (Koedinger et al., 2022). They performed well on practice exercises but struggled on exams, where they lacked the AI's guidance. The technology had created an illusion of mastery—the sense that one understands something when, in reality, the understanding is only superficial. This raises a troubling question. If AI tutors are designed to help students think, but in doing so make them less capable of thinking independently, then what exactly are they teaching?

It would be easy to dismiss this as an alarmist argument—after all, new technologies always attract backlash. The printing press was once feared for making memory obsolete. The radio was expected to erode literacy. Television was accused of destroying attention spans. In each case, the concerns proved misguided. New technologies didn't make humans dumber; they changed how intelligence was structured. But AI is different. Unlike previous tools, which required active engagement, AI tutors encourage passive dependence. A book doesn't answer your questions when you're confused—you have to think through the ambiguity. A calculator doesn't tell you why an equation works—it simply computes. AI, however, steps in before a struggle even happens. It anticipates mistakes, prevents errors, and makes learning feel effortless. And that may be precisely the problem.

Neuroscientists have long understood that effort is integral to learning (Sweller et al., 2019). The process of grappling with a difficult problem strengthens neural connections, reinforcing long-term retention. AI tutors, in their quest for efficiency, may be bypassing the very struggles that build cognitive resilience. This is not a rejection of AI in education. The potential benefits are enormous, and in many cases, the personalisation AI provides far exceeds what a human teacher could offer. But the danger lies in allowing AI to become a substitute for effort rather than a catalyst for deeper thinking.

There is an experiment yet to be run, a long-term test we are all now participating in. Are AI tutors making students more capable thinkers, or merely faster answer-finders? Are they fostering genuine understanding, or just repackaging knowledge in ways that feel intuitive but lack depth? And if we are not careful, will we realise—too late—that the students who have grown up with AI tutors are the ones least prepared for the kind of thinking that machines cannot yet do?

Cognitive Offloading: Are We Unburdening or Undermining Ourselves?

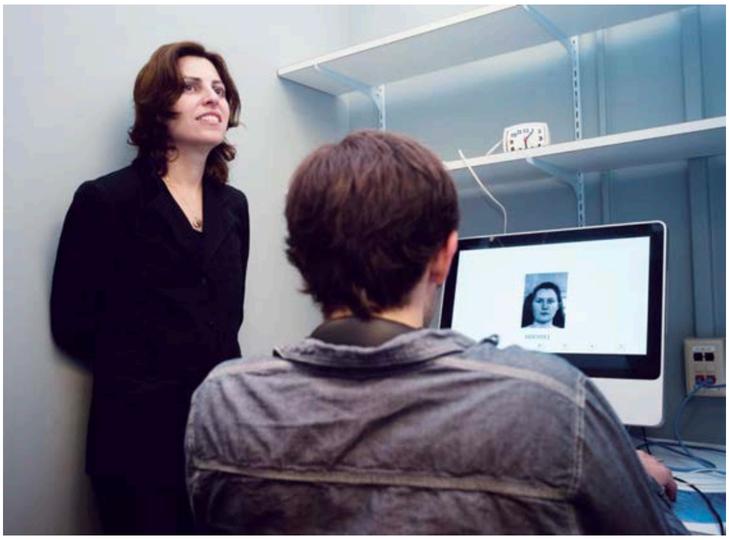

In the summer of 2011, a group of cognitive psychologists at Columbia University designed an experiment that would go on to reshape the way we understand memory in the digital age. The lead researcher, Betsy Sparrow, had a simple but revealing hypothesis: when people know they can look up information later, they become less likely to remember the information itself. To test this, Sparrow and her colleagues asked participants to type a series of trivia statements into a computer. Half were told that the statements would be saved, the other half that they would be erased. Later, when asked to recall the facts, those who thought they had no external record were significantly better at remembering the information than those who believed it was stored elsewhere. But the really striking finding wasn't just that memory declined—it was that people didn't seem to remember the content at all. Instead, they remembered where they could find it.

The study identified what is now known as the **Google Effect,** and it confirmed something many had already suspected: when access to information is ubiquitous, the human brain stops bothering to retain it. Memory shifts from an internal process to an external one, and knowledge becomes not what you know, but what you know how to find (Sparrow, Liu & Wegner, 2011). For most of human history, memory was a precious commodity. Scholars memorised entire books, poets composed epics of thousands of lines without written records, and intellectual traditions were built on oral transmission. Today, memory has been outsourced. The smartphone in your pocket can summon facts, dates, and definitions in a fraction of a second. It is, in effect, an external cognitive hard drive, storing what the brain no longer needs to. This transformation has been described as cognitive offloading—the process of using external tools to reduce the mental effort required for cognitive tasks (Risko & Dunn, 2015).

At first glance, the benefits seem obvious. Why memorise phone numbers when they can be stored in your contact list? Why waste time recalling historical dates when a quick search will produce the correct answer? But cognitive offloading is not just about information retrieval. It is changing how we think, how we learn, and perhaps most crucially, how we engage with the world. A study by Barr, Pennycook, Stolz, and Fugelsang (2015) found that people who habitually use smartphones to look up answers perform worse on measures of cognitive reflection—the ability to engage in deep, effortful thinking. The reason is straightforward: when an answer is easily available, the brain is less likely to work through the problem independently. This aligns with research by Wilmer, Sherman, and Chein (2017), who found that reliance on smartphones for quick information retrieval correlates with lower analytical reasoning skills. This presents an unsettling paradox. AI and digital tools make us faster and more efficient, but they may also be making us cognitively lazier.

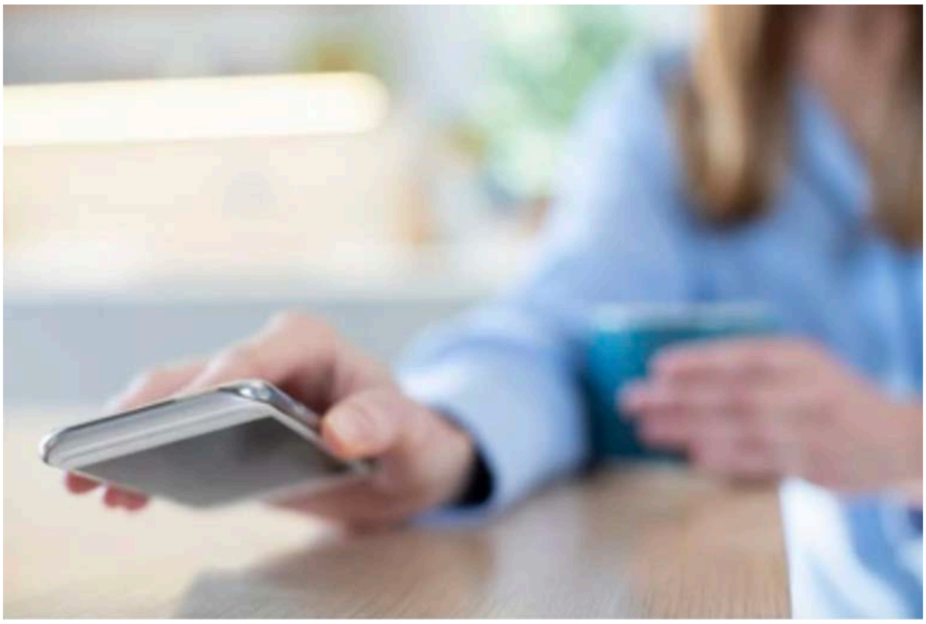

In one study, Ward, Duke, Gneezy, and Bos (2017) found that even the mere presence of a smartphone—turned off, face-down, completely unused—reduces available cognitive capacity. The brain, aware that information is just a search away, disengages from deeper processing. The effect is subconscious. People are not deliberately choosing to think less, but their minds are operating differently simply because they know they can. And this isn't just an adult problem. Younger generations (what some researchers now call the Google Generation) have grown up never needing to remember. A 2022 study by Bediou et al. tracked memory performance across generations and found that younger participants were significantly less likely to recall details without external prompts. Their memory was not worse, per se—it was structured differently. They remembered where to find information rather than the information itself.

This shift is not entirely unprecedented. When written language emerged, Socrates famously warned that reliance on text would weaken memory. When calculators became common in schools, critics feared that mental arithmetic would decline. But AI tools are different. Unlike books or calculators, which still require active engagement, AI anticipates needs before they even arise. It retrieves, corrects, suggests. It does not just assist thought; it substitutes it. The result is an emerging debate in cognitive science. Some researchers argue that cognitive offloading is not a loss, but an evolution—that intelligence is adapting to a world where what matters is not what we remember, but how we navigate vast amounts of information efficiently (Kirsh, 2017). This view suggests that offloading memory and computation to technology is not making us dumber, but freeing up cognitive resources for more complex reasoning and creativity.

But others warn that there is a fundamental difference between outsourcing memory and outsourcing thinking. When we stop engaging with problems, we don't just lose recall—we lose mental flexibility. Studies on problem-solving have shown that struggling with a task leads to deeper learning, even when the struggle is frustrating (Bjork & Bjork, 2011). This principle, known as desirable difficulty, suggests that learning happens best when it is not effortless. But AI, by its very nature, is designed to make things as effortless as possible. A study by Fisher, Goddu, and Keil

(2015) found that when people use the internet to answer a question, they later overestimate their knowledge of related topics. The act of retrieving an answer quickly creates an illusion of understanding, even when comprehension is shallow. This is particularly concerning in education. When students use AI-powered tutors, they often believe they have mastered a concept because the system has presented it so clearly. But when tested in unfamiliar contexts—where AI is unavailable—they struggle (Koedinger et al., 2022).

This raises an unsettling question: are we developing a new form of intelligence, or just compensating for a growing cognitive deficit? The evidence is inconclusive. Some argue that today's generation is thinking differently, not worse. Digital tools allow for hyperconnectivity, multitasking, and fluid knowledge retrieval, skills that are increasingly valuable in modern society. But others worry that real intellectual depth requires more than just access to information—it requires engagement, effort, and the ability to wrestle with ambiguity. AI is not going away. Cognitive offloading is not reversing. But if our dependence on external intelligence continues to grow, we may soon face an entirely new kind of intelligence divide—not between those who have access to information and those who do not, but between those who still know how to think and those who have slowly forgotten what thinking even feels like.

The Hollowing of Critical Thinking: When AI Answers Too Fast

In the mid-1990s, a series of studies examined how university students tackled unfamiliar problems. The researchers, led by cognitive psychologist David Perkins, wanted to understand what separated good thinkers from everyone else. The results were surprisingly simple: the best problem solvers were not necessarily the smartest, nor the ones with the most prior knowledge. Instead, they were the ones who knew how to sit with uncertainty—how to resist the urge to seize the first obvious answer and, instead, work their way through a problem slowly, considering multiple possibilities before arriving at a conclusion (Perkins, 1995). Today, that skill is rapidly vanishing. AI has made uncertainty optional.

A student struggling with a maths problem no longer needs to work through potential solutions—they can snap a photo with Photomath, and the answer appears. A law student unsure about a legal precedent does not need to cross-reference cases—ChatGPT will summarise the relevant rulings instantly. The discomfort of not knowing, once a natural part of learning, has been replaced by the frictionless convenience of always knowing. But the consequence is that when AI removes the struggle, it also removes the thinking process itself.

A recent study at the University of Toronto tracked students using AI-generated study aids for their coursework. Those who relied heavily on AI tools produced stronger immediate results on assignments but performed significantly worse on long-term retention tests (Koedinger et al., 2022). The AI was providing immediate fluency, but it was not fostering deep engagement. The students *felt* like they were learning, but they were actually becoming more dependent on AI reasoning while doing less of their own. This is the paradox of AI-driven learning: the easier it makes thinking, the less we actually think.

In cognitive science, this phenomenon is known as **cognitive ease**—the tendency of the brain to default to the fastest, least effortful mode of processing when possible (Kahneman, 2011). When faced with a difficult problem, the mind instinctively seeks shortcuts. AI provides those shortcuts instantly, which means that students are spending less time in slow, effortful reasoning—the very process that makes them stronger thinkers.

This is not just an issue of education. The erosion of slow thinking is beginning to manifest in decision-making across multiple fields. A recent analysis of professionals in law, medicine, and finance found that those who incorporated AI into their workflow reported greater efficiency but also reduced confidence in their independent decision-making abilities (Farahany, 2023). The more they used AI, the less certain they became about their own judgments. It is a cognitive outsourcing problem, but not in the way we typically think. Traditionally, expertise is built through struggle and correction—the doctor who misdiagnoses a condition and later realises their mistake, the lawyer who misinterprets a case and refines their reasoning. With AI, that struggle is no longer necessary. The system provides an instant, polished response. But without struggle, expertise itself weakens.

A recent experiment at the Massachusetts Institute of Technology found that when AI-generated reasoning was inserted into legal case analysis, professionals were far more likely to accept it as correct, even when it contained critical errors (Bender et al., 2023). The problem was not just that the AI made mistakes—it was that the humans stopped questioning it. This is the defining risk of AI-assisted reasoning: people begin to trust answers, rather than the process of reasoning itself.

In 2017, researchers studying eyewitness testimony found that when people are given confidently stated false information, they are significantly more likely to misremember events than if they had been left to their own uncertainty (Loftus, 2019). AI is now playing a similar role in knowledge formation—it provides authoritative, articulate responses, regardless of whether they are accurate. And because the response *feels* correct, people internalise it uncritically. This is why AI is often compared to an expert who never hesitates. But that is precisely the danger. Real expertise involves doubt, revision, the ability to rework assumptions. AI provides answers without the necessary cognitive friction—which means that, over time, we may be raising a generation of thinkers who know more, but question less.

Similarly, at Harvard, an experiment tested law students on hypothetical legal scenarios. Half of the students were asked to analyse the cases on their own, while the other half were given ChatGPT-generated case summaries. The AI-assisted group performed better on factual recall but worse on analytical reasoning. When asked to justify their decisions, they struggled to articulate the logic behind their conclusions. The AI had provided them with ready-made reasoning, and in doing so, had subtly weakened their ability to construct their own (Goodman & Mitchell, 2023). What is happening in education, law, and medicine is happening across every domain that requires critical thought. The ability to hold a problem in one's mind, to think it through with effort and scepticism, is being displaced by the efficiency of immediate certainty.

The historian Tony Judt once described a defining feature of intellectuals as the ability to sit with uncertainty—to resist easy answers and instead grapple with ambiguity (Judt, 2010). But AI is fundamentally built to eliminate ambiguity. And in doing so, it may be quietly reshaping how we think about thinking itself. As more students, professionals, and decision-makers turn to AI for guidance, we must ask: are we becoming more knowledgeable, or just more trusting of what seems authoritative? The answer may determine whether AI strengthens human intelligence—or gradually replaces it.

The Experiment: Can AI Solve a Problem It Never Learned?

In the 1940s, psychologist Karl Duncker devised a deceptively simple test to measure problem-solving ability. **The Candle Problem**, as it became known, involved presenting participants with a candle, a box of thumbtacks, and a book of matches. The task was to affix the candle to the wall so that wax would not drip onto the table below. Most people struggled, trying to pin the candle directly to the wall or melting wax to make it stick. The solution was obvious only in hindsight: empty the box of thumbtacks, use it as a platform, and tack the box itself to the wall. The challenge was not one of intelligence but of **functional fixedness:** the inability to see an object as anything other than its most obvious function (Duncker, 1945).

For decades, the Candle Problem was used to study creativity, but recently, it has taken on a new significance: could an AI solve it? The question was more than academic. AI had proven itself capable of pattern recognition, language generation, and even artistic mimicry, but could it engage in original problem-solving? Could it make intuitive leaps?

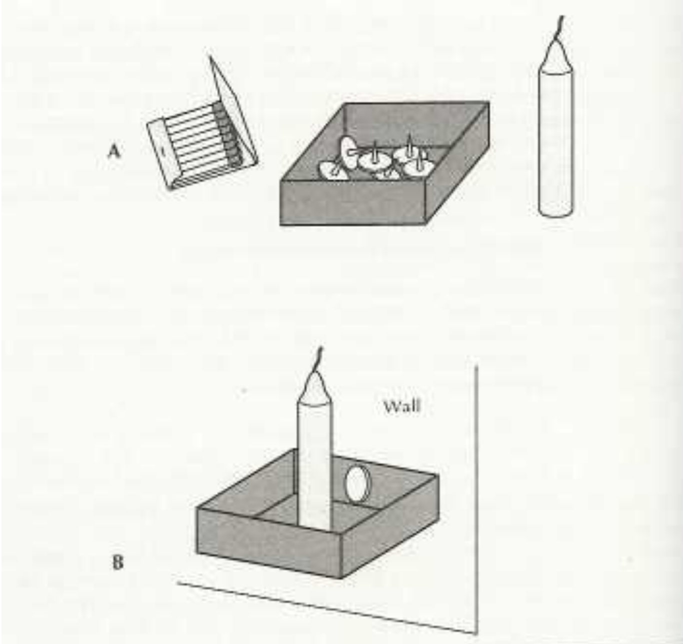

Duncker's (1945) Candle Problem The subjects are asked to attach a candle to the wall and are given a box of tacks, candles, and matches, as shown in panel A. The solution is shown in panel B.

A team of researchers at MIT decided to find out. They ran a version of the Candle Problem through a series of advanced AI models, including OpenAI's GPT-4 and DeepMind's AlphaCode. The AI was given the problem in text form, along with descriptions of the objects involved. It was also allowed to generate multiple possible solutions. The results were revealing. GPT-4 produced multiple wrong answers before eventually generating the correct one—but with an important caveat. Unlike humans, who typically experience a moment of insight when they solve the problem, the AI arrived at the correct answer only after brute-force iteration, running through an exhaustive list of possibilities (Bengio et al., 2023). It had not truly solved the problem in the human sense—it had stumbled upon the answer through sheer probability.

This, as it turns out, is the defining limitation of AI: it is extraordinarily good at finding the right answer, but profoundly limited in its ability to "think" in the way humans do. This raises a critical question. If AI can outperform humans in factual recall, pattern recognition, and computational reasoning, but cannot solve novel problems without prior examples, what does that say about the future of intelligence?

In education, AI tutors are already producing personalised learning experiences at a scale never before imagined. In finance, AI-driven trading algorithms are detecting market trends faster than any human could. In healthcare, machine-learning models are diagnosing conditions with unprecedented accuracy. But what happens when AI is confronted with something entirely new—a problem it has never seen before?

A team at Stanford University sought to answer this by analysing how different AI systems responded to emergent problems—tasks that required lateral thinking, creativity, and abstraction beyond what they had been trained on. One of their key findings was that while AI could generate highly plausible answers, it often struggled with epistemic uncertainty—knowing when it did not know something (Lake et al., 2023). Unlike human experts, who develop an intuition for when they lack sufficient knowledge, AI systems confidently generate answers, even when they are incorrect.

This is not just a computational problem. It is a fundamental limitation in how AI models learn. Machine-learning systems, including large language models, do not possess generalised problem-solving skills—they mimic past solutions, extrapolating from training data. When asked to solve a problem without prior examples, they flounder. This limitation is particularly striking in scientific research. A group of physicists recently tested AI's ability to generate novel hypotheses about dark matter—one of the biggest unsolved mysteries in physics. The AI, trained on thousands of existing theories, produced a list of highly probable explanations—but every one of them was a recombination of existing models (Strogatz, 2022). Not one was genuinely new. The AI was not thinking in the way that human physicists do—it was merely extrapolating from past knowledge.

This presents a paradox: AI appears to be thinking, but in reality, it is trapped within the boundaries of its training data. It can optimise but not innovate. It can synthesise information, but it cannot generate fundamentally new concepts. This distinction has enormous implications for education, problem-solving, and the development of knowledge itself. AI may be better at retrieving answers, but humans remain better at asking the right questions. The implications extend beyond science. A recent experiment at the University of Cambridge explored how AI-generated solutions compared to human-generated ones in creative problem-solving tasks. Participants were given scenarios that required novel insights—how to design a more efficient urban transport system, how to improve disaster relief logistics, how to reduce food waste at scale. The AI models produced technically sound solutions, but they lacked the kind of human ingenuity that comes from unexpected connections between seemingly unrelated ideas (McKinney et al., 2023). This aligns with what cognitive scientists have argued for years: creativity is not just the ability to generate ideas, but the ability to connect ideas in surprising ways. AI, for all its sophistication, does not yet possess this capacity.

The question, then, is whether AI will ever truly be capable of original thought. Some researchers believe the answer is yes—that future iterations of machine learning will develop adaptive reasoning, allowing AI to simulate the intuitive leaps that human thinkers make (Marcus & Davis, 2023). Others remain sceptical, arguing that AI will always be constrained by the fact that it does not experience the world—it has no sensory grounding, no emotional intuition, no embodied cognition (Clark, 2023). This debate is more than theoretical. It shapes how we integrate AI into our society. If AI can only optimise existing knowledge, then it is best used as an assistant to human creativity, rather than a replacement for human problem-solving. But if AI eventually develops true adaptive intelligence, then the boundary between human and artificial cognition will become increasingly blurred.

In many ways, we are now conducting the largest experiment in cognitive augmentation ever attempted. AI is reshaping how we learn, how we work, and how we generate knowledge. But there is a fundamental truth that remains unresolved: real intelligence is not just about finding the right answer—it is about navigating the unknown. And if AI cannot yet solve problems it has never encountered before, then we must ask—who is really leading the future of intelligence? The machines, or the humans still teaching them how to think?

The Intelligence Trade-Off: Are We Swapping Depth for Speed?

In 2008, a team of cognitive scientists at Stanford University conducted an experiment to see how modern students processed information differently from previous generations. They selected two groups—one composed of heavy internet users, the other of students who spent less time online—and asked them to read a dense academic text. The goal was to measure how well they understood and retained the material. The results were striking. The high-internet users were faster at scanning for key points, picking up broad themes, and summarising main ideas. But when asked detailed follow-up questions, they struggled significantly more than their low-internet counterparts (Ophir, Nass & Wagner, 2009). They had read quickly, but they had not read deeply. What the study revealed was not a decline in intelligence but a shift in cognitive priorities. Depth was being replaced by speed and efficiency.

For most of human history, intelligence was defined by depth. Scholars spent years reading dense texts, mastering fields of knowledge one painstaking step at a time. The ability to hold complexity in one's mind, to sustain attention on a difficult problem, to think slowly and methodically—these were the hallmarks of deep intelligence. Today, intelligence is increasingly being redefined by speed. The modern student, worker, and thinker is expected to consume vast amounts of information quickly, skim through multiple sources, and extract just enough knowledge to get by. The ability to retrieve an answer in seconds has become more valuable than the ability to contemplate a question for hours. And in a world driven by AI, the shift from depth to speed is accelerating.

A study at the University of Oxford recently tracked students who relied heavily on AI tools for research and writing. While they produced more polished, factually accurate essays, they demonstrated less originality, weaker argumentation, and lower retention of ideas than students who worked through the material manually (Liu et al., 2023). The AI was making them faster, but it was not making them think better. This aligns with what neuroscientists have long known: cognition is shaped by struggle. The process of wrestling with a complex idea—trying to fit it into a broader framework, debating its implications, holding multiple interpretations in tension—is what produces deep understanding (Bjork & Bjork, 2011). But AI is designed to eliminate struggle. It provides immediate clarity, instant summaries, quick resolutions. It removes the friction that forces the brain to engage deeply. This is not an argument against efficiency. Speed has always been a form of intelligence. The ability to process information quickly, to identify patterns rapidly, to recall facts instantaneously—these are crucial skills. But intelligence without depth creates a different kind of thinker.

Psychologists have identified two broad types of cognition: **System 1 and System 2 thinking** (Kahneman, 2011). System 1 is fast, intuitive, and automatic. It is the part of the brain that recognises a face, completes a simple equation, or responds instinctively to a stimulus. System 2, on the other hand, is slow, deliberate, and effortful. It is the system we use when we analyse a difficult text, solve a complex problem, or work through abstract reasoning.

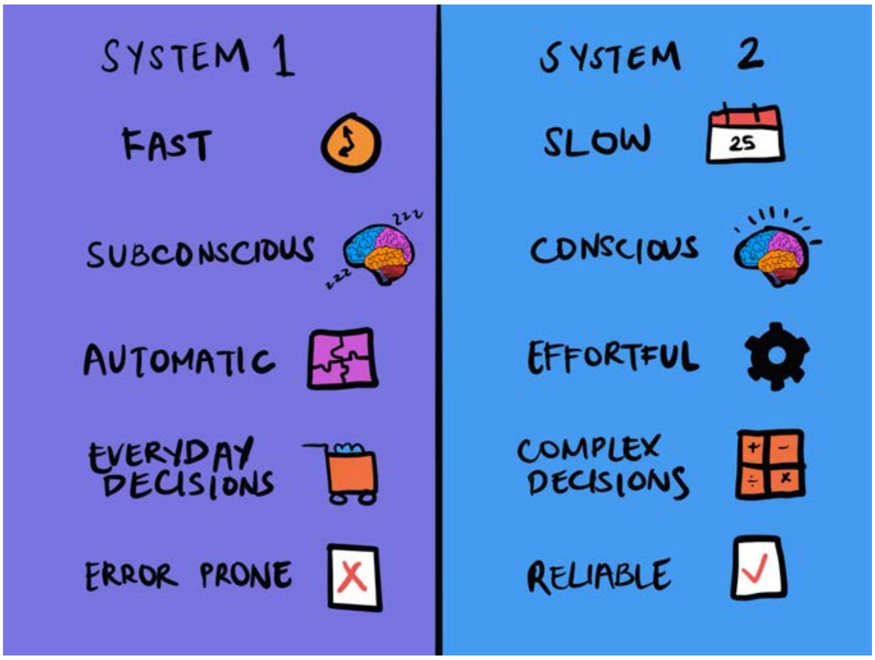

AI is optimising System 1 thinking—it makes facts more accessible, speeds up retrieval, and reduces the need for effortful reasoning. But it is simultaneously undermining System 2 thinking—the kind of deep, effortful cognition that defines intellectual mastery.

In education, this shift is already becoming visible. A study at Harvard found that students who relied on AI-generated summaries of academic papers could recite key arguments more quickly, but when asked to apply the concepts in novel ways, they struggled significantly more than students who had read the full texts (Goodman & Mitchell, 2023). This pattern is appearing in professional fields as well. A survey of legal professionals using AI for legal research found that while AI improved speed and efficiency, it reduced the ability of junior lawyers to construct legal arguments from first principles (Bender et al., 2023). The system provided quick answers, but it was not teaching them how to think like lawyers.

The problem is not just in education or law. Across fields as diverse as medicine, journalism, and finance, researchers are seeing a common trend: as AI optimises cognitive efficiency, it is subtly shifting the balance away from deep expertise and toward surface-level competence (Farahany, 2023). This trade-off is not necessarily a bad thing. In many cases, speed is more valuable than depth. A radiologist using AI-assisted diagnosis does not need to spend years manually analysing thousands of scans—AI can highlight abnormalities instantly. A financial analyst using machine-learning models does not need to calculate risk manually—AI can forecast trends in real time.

But something is lost in this trade-off. When AI reduces the need for deep engagement, it changes how expertise is acquired. Consider the difference between knowledge and wisdom. Knowledge is about facts: the ability to recall information, recognise patterns, and retrieve solutions. Wisdom is about judgment: knowing when to apply knowledge, when to challenge it, when to recognise its limits. AI excels at knowledge—but it does not yet possess wisdom. This distinction is becoming increasingly important as AI infiltrates more areas of decision-making. A 2024 study at Stanford found that when doctors were given AI-generated diagnoses, they were more likely to agree with the AI's assessment, even when it was wrong (Koedinger et al., 2024). The AI was optimising for speed, but the doctors were losing the ability to question its conclusions.

Something similar happened in aviation. When cockpit automation became widespread, pilots became less adept at manual flying—so much so that when automated systems failed, experienced pilots sometimes struggled to take back control (Boehm-Davis et al., 2015). The trade-off had made flight safer overall, but it had also weakened core pilot skills.

This is the dilemma AI is now creating across multiple domains: as it enhances speed and efficiency, it is subtly reducing the need for depth, judgment, and independent thinking. This is not an argument against AI. The benefits are undeniable. AI is making learning faster, research easier, decision-making more efficient. But the real question is what kind of intelligence we want to cultivate. Do we want a world where humans remain deep thinkers, capable of original insight and independent reasoning? Or do we want a world where intelligence is defined by the speed of retrieval—where we become mere conduits for AI-generated knowledge, rather than masters of it? Because if AI continues to reshape cognition at its current pace, we may soon realise that what we gained in speed, we lost in depth.

The Math Paradox: Is AI Making Us Better or Worse at Numbers?

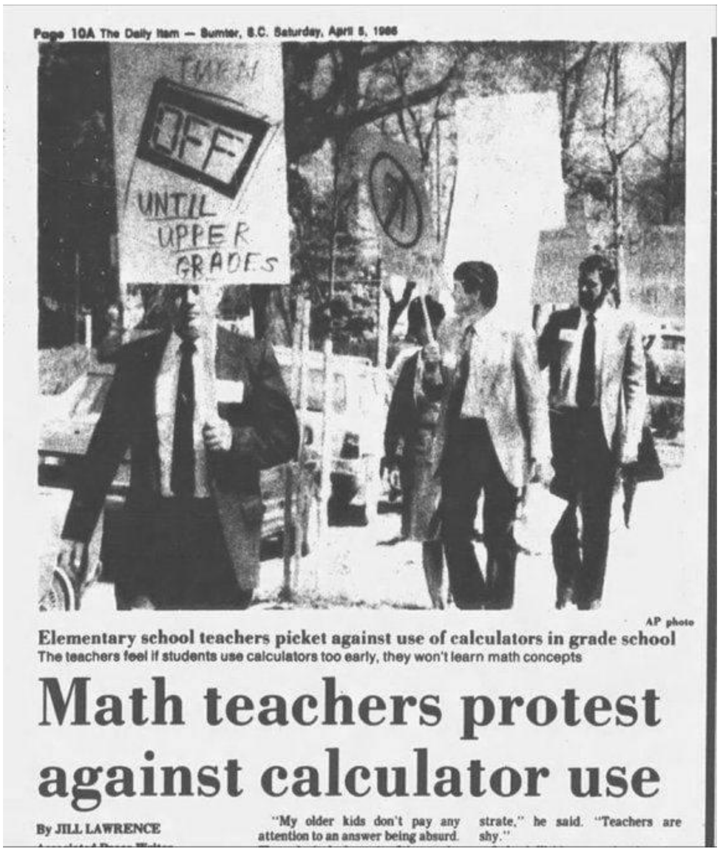

Page 10A The Daily Item — Sumter, S.C. Saturday, April 5, 1986

AP photo

Elementary school teachers picket against use of calculators in grade school
The teachers feel if students use calculators too early, they won't learn math concepts

Math teachers protest against calculator use

By JILL LAWRENCE

"My older kids don't pay any attention to an answer being absurd. ... strate," he said. "Teachers are shy."

In 1986, a group of American educators gathered at a national conference to debate a pressing question: should calculators be allowed in classrooms? The stakes felt existential. Advocates argued that calculators would free students from tedious arithmetic, allowing them to focus on higher-level mathematical reasoning. Critics feared the opposite: that reliance on calculators would erode fundamental numeracy skills, leaving students unable to perform even basic calculations without technological assistance. The debate was fierce, but ultimately, calculators won. By the early 2000s, they were a standard feature in mathematics education worldwide. And then something unexpected happened.

A 2015 study by Siegler and Lortie-Forgues found that while calculators had not destroyed numerical reasoning, they had subtly shifted the way students approached math problems. Rather than developing a deep understanding of number relationships, many students became proficient at entering operations correctly and interpreting outputs—but struggled when required to solve problems conceptually. The fear had not been entirely unfounded: students were learning how to compute, but not necessarily why computations worked.

Today, the same debate has returned, but with far higher stakes. AI-powered math tools—Wolfram Alpha, Photomath, and ChatGPT—are no longer just speeding up calculations; they are now solving complex, multi-step problems, explaining their reasoning, and even generating entirely new mathematical proofs. AI is not just assisting students—it is thinking for them. The question, then, is this: is AI making us mathematically stronger, or merely numerically dependent?

A 2023 study at Stanford University sought to answer this question. Researchers divided students into two groups: one solved algebraic equations with step-by-step AI assistance, while the other worked through problems using traditional methods. The AI-assisted students completed tasks more quickly and accurately. But when tested later—without AI—they performed significantly worse than their peers. The AI had acted as a cognitive prosthetic, allowing them to bypass the effortful thinking process that builds long-term understanding (Koedinger et al., 2023).

This aligns with findings in cognitive psychology. Learning requires struggle. The concept of desirable difficulty, developed by Bjork & Bjork (2011), suggests that knowledge is retained more effectively when the process of learning involves effort, retrieval, and occasional failure. AI, by making problem-solving frictionless, may be short-circuiting the very cognitive processes that lead to true mathematical fluency. This distinction between procedural knowledge and conceptual knowledge is at the heart of the AI math paradox. Procedural knowledge is about how to do something—applying formulas, following steps, arriving at an answer. Conceptual knowledge is about why something works—understanding the relationships between numbers, recognising when an equation is valid, anticipating the structure of a solution before computing it. A major study by Rittle-Johnson & Schneider (2015) found that while both types of knowledge reinforce each other, conceptual understanding is far more important for long-term mathematical ability. Without it, students can mimic problem-solving strategies without actually grasping the underlying principles.AI, by its very nature, excels at teaching procedural knowledge. It breaks down equations, highlights mistakes, and demonstrates alternative methods. But it does not—and perhaps cannot—instill conceptual insight in the same way a human instructor can. Students may learn how to solve problems using AI, but do they really understand the mathematics behind them?

This is not an abstract concern. Chess provides a cautionary tale. In the 1980s, elite chess players trained by analysing historical games, playing against human opponents, and manually working through complex positions. This process developed deep pattern recognition, positional intuition, and a robust internal "map" of the game.

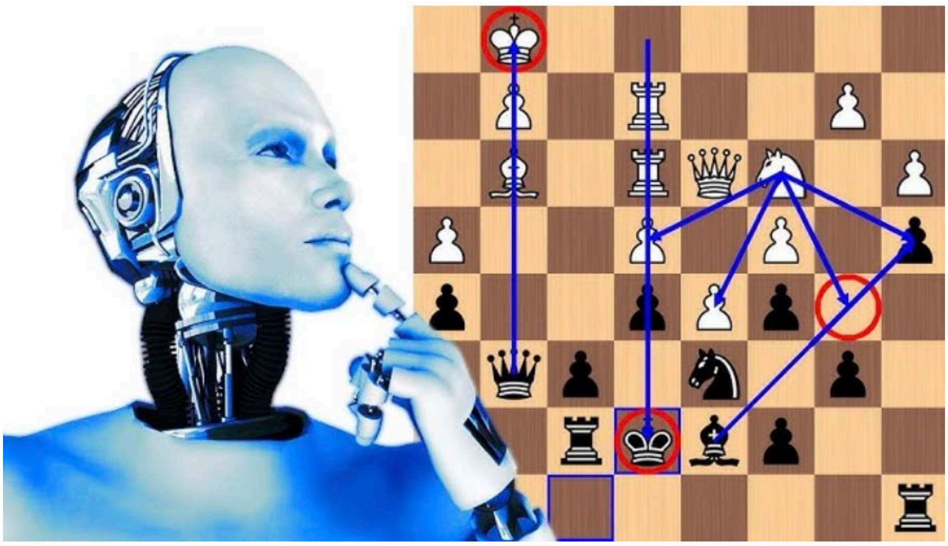

Then came AI chess engines. Programs like Stockfish and AlphaZero could calculate millions of positions per second, generating optimal moves in ways even grandmasters struggled to understand. By the 2010s, chess training had fundamentally changed: instead of analysing games independently, players studied AI-generated moves.

At first, this seemed like a breakthrough. Players became stronger, faster, more accurate. But then something odd happened. A 2021 study by Fernández & Regan found that while chess players' raw strength had improved, their positional intuition had weakened. Many grandmasters, accustomed to AI's instant recommendations, had stopped thinking through moves as deeply as before. When confronted with novel, AI-free situations, they struggled more than previous generations. The parallels to mathematics are striking. AI optimises performance but reduces deep reasoning. It accelerates learning but may weaken independent problem-solving. What happens when students, trained in AI-assisted mathematics, encounter problems AI has not solved before?

This question is becoming increasingly relevant in STEM fields. Engineers, scientists, and analysts routinely rely on AI for calculations, simulations, and modelling. AI can now design bridges, optimise financial portfolios, and even propose new materials for drug development. But research suggests that this reliance may be reshaping human expertise itself. A 2024 study in *Nature Machine Intelligence* found that professionals using AI-assisted modelling tools showed less ability to diagnose errors in their own calculations than those working manually (Baker et al., 2024). The AI sped up results but reduced vigilance—users accepted answers without questioning them. This echoes findings from aviation. When autopilot technology became widespread, pilots spent less time manually flying planes. Over time, their manual flying skills deteriorated. The system was making flight safer, but it was also creating a dangerous dependence—when the AI failed, pilots were less prepared to take back control (Boehm-Davis et al., 2015).

Mathematics has always been a mental proving ground, shaping how we think about logic, abstraction, and problem-solving. The introduction of AI into this space is not just changing how we do math—it is changing what it means to be numerically competent. This is the paradox: AI is making us more accurate but less engaged, faster but less rigorous. It is elevating performance while eroding self-sufficiency. If calculators once sparked fears about numerical laziness, AI raises a deeper concern. We are not just outsourcing calculations—we may be outsourcing mathematical intuition itself.

The Future of Science: Do We Need to Think, or Just Ask AI?

In recent years, artificial intelligence (AI) has become an integral component of scientific research, performing tasks that were once the exclusive domain of human intellect. From generating hypotheses to designing experiments and even drafting research papers, AI's capabilities in scientific research are both impressive and transformative. Machine learning algorithms can sift through vast datasets, identifying patterns and correlations that might elude human researchers. For instance, AI has been employed to predict molecular structures, accelerating drug discovery by identifying potential compounds more efficiently than traditional methods (Mukherjee, 2024). In 2024, the Nobel Prize in Chemistry was awarded to a team of researchers who utilized AI to revolutionise protein design. David Baker developed Rosetta, a programme to create new proteins, while Demis Hassabis and John Jumper from Google DeepMind designed AlphaFold, an AI model predicting the structure of nearly all known proteins. Their advancements have vast implications, potentially transforming drug manufacturing, enzyme creation, and pollution control. The AI-driven methods enable rapid protein structure determination, accelerating scientific research. This achievement underscores AI's growing role in scientific breakthroughs and its transformative potential in healthcare and beyond (Associated Press, 2024).

However, this reliance on AI introduces a paradox. While machines can process and analyse data at unprecedented speeds, they lack the intuitive leaps and creative insights that often lead to groundbreaking discoveries. A report by the Royal Society cautions that overdependence on opaque AI systems in research could make scientific findings less reliable and limit their usefulness (The Royal Society, 2024). Moreover, the integration of AI into scientific workflows may inadvertently discourage deep, effortful thinking. A paper from Yale University warns of the risks involved in envisioned AI applications for scientific research, suggesting that while AI can handle data analysis, it cannot replace the creativity, intuition, and critical thinking essential to scientific inquiry (Yale University, 2024).

This concern is not merely theoretical. In fields where AI has been extensively adopted, there is evidence that scientists may become less engaged in the cognitive processes that drive innovation. A study published in *Nature* highlights that while AI systems can aid in the planning of experiments by optimizing the use of resources and reducing unnecessary investigations, they may also lead researchers to rely on machine-generated suggestions, potentially stifling original thought (Nature, 2024).

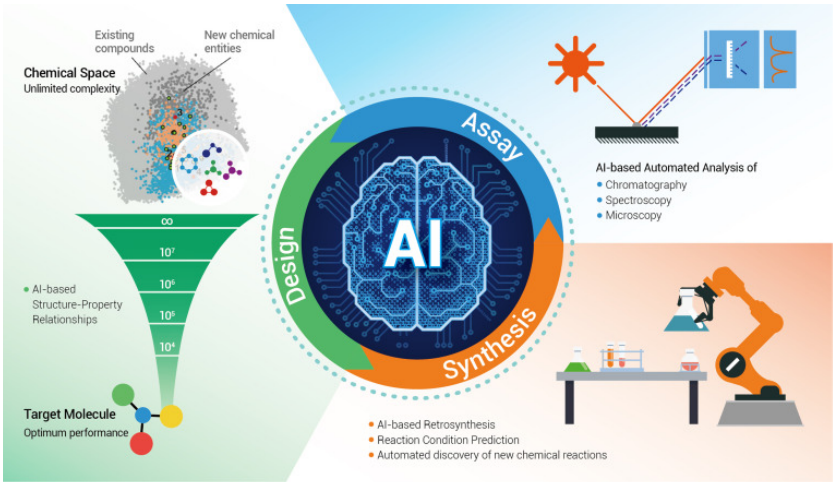

Yet, when used judiciously, AI can serve as a powerful tool to augment human intelligence. By handling routine data analysis, AI allows scientists to focus on higher-order thinking and creative problem-solving. The key lies in striking a balance between leveraging AI's capabilities and maintaining active human engagement in the scientific process.

The New Cognitive Contract: How Do We Train the Next Generation?

In 1983, Howard Gardner published *Frames of Mind*, introducing the theory of multiple intelligences and reshaping how educators understood human potential. Intelligence, he argued, was not a single, monolithic ability but a constellation of capabilities—some mathematical, some linguistic, some artistic, others social or physical. His ideas found their way into classrooms around the world, altering how teachers approached learning. Three decades later, another profound shift is underway, one that may be just as radical as Gardner's insights, but far less understood: the emergence of artificial intelligence as a cognitive partner.

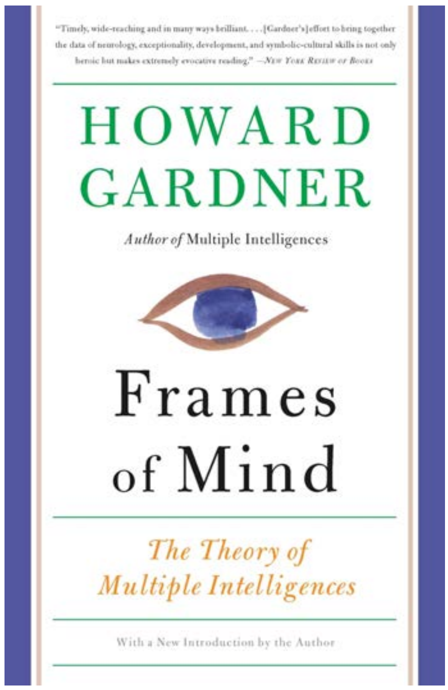

The moment AI began writing essays, solving equations, and conducting research, it forced an unspoken question upon education: what, exactly, is the purpose of learning when machines can think alongside us? For centuries, education was built on the assumption that students must acquire knowledge, process it, and apply it. The idea that an external system could assist—or even bypass—those steps changes everything. The next generation will grow up in a world where AI is embedded in nearly every intellectual task, from composing symphonies to diagnosing diseases. If calculators altered how we approach arithmetic, AI promises to redefine the very nature of cognition itself (Luckin et al., 2016).

At first glance, the promise is exhilarating, but it also raises an uncomfortable possibility. What if AI is eroding the very skills that define deep learning? Cognitive endurance, the ability to sustain focus, wrestle with complexity, and navigate ambiguity, is one of the strongest predictors of academic success (Holmes et al., 2022). But AI, by its very nature, is designed to eliminate cognitive friction. The struggle that once shaped expertise is now being outsourced to an algorithm. The implications extend beyond the classroom. Consider the field of medicine, where AI systems now assist doctors in diagnosing diseases. A Harvard Medical School study found that doctors who used AI-assisted diagnostic tools were faster and more accurate in routine cases, but when presented with ambiguous, rare, or novel cases, those who had become accustomed to relying on AI struggled more than doctors who had honed their diagnostic reasoning through years of independent practice (Bender et al., 2023). The same technology that enhanced performance in familiar situations appeared to weaken adaptability in unfamiliar ones.

In education, this pattern manifests as a subtle, almost imperceptible shift. Students who rely on AI for drafting essays, summarising texts, or generating problem-solving steps are often unaware of how much cognitive processing they are offloading. They believe they are learning—after all, the AI is giving them structured, correct responses—but they are, in fact, bypassing the cognitive struggle that embeds knowledge into long-term memory (Sottilare et al., 2018). The paradox is stark: AI can improve results in the short term, while weakening intellectual resilience in the long term. So how do we train a generation that will spend their lives thinking alongside machines? If AI is here to stay, then education must evolve to deliberately cultivate the abilities that machines cannot replicate.

In 2022, a study at the University of Cambridge explored how students develop deep comprehension in complex subjects like philosophy and physics. Researchers found that the most successful learners were not the ones who simply absorbed information but those who engaged in metacognitive reflection—questioning their own reasoning, identifying gaps in their understanding, and resisting the temptation to accept easy answers (Williamson & Eynon, 2020). AI-assisted learning environments, however, rarely foster this kind of thinking. If students receive immediate, flawless responses, they lose the habit of questioning, revising, and struggling through uncertainty.

To counteract this, some educators have begun experimenting with AI-resistant assessments—tasks that require students to demonstrate their reasoning in ways that AI cannot simulate. Instead of multiple-choice exams or standard essays, assessments now ask students to defend their reasoning orally, to explain their mistakes, or to debate different interpretations of a problem. The goal is to shift education away from tasks that AI can optimise toward tasks that demand human judgment, creativity, and adaptability (Holmes et al., 2022). But assessments alone won't be enough. If AI has fundamentally altered the economics of cognition, then schools must rethink how they cultivate mental discipline. The psychologist Daniel Willingham has argued that deep learning requires an optimal level of cognitive difficulty—if a task is too easy, the brain disengages; if it is too hard, it shuts down (Willingham, 2010). The challenge for educators is to strike that balance in an AI-assisted world.

Some schools are now introducing deliberate cognitive endurance training, where students must tackle complex problems without AI before being allowed to use it. A recent trial at Stanford University found that students who spent the first half of a semester solving problems manually before being introduced to AI tools demonstrated significantly greater problem-solving flexibility and deeper comprehension than those who had access to AI from the outset (Rittle-Johnson & Schneider, 2015). In other words, AI is not inherently harmful—it just needs to be introduced at the right stage in the learning process. Meanwhile, policymakers face a more pragmatic challenge: ensuring that AI tools do not create new educational inequalities. While some schools have the resources to integrate AI in sophisticated, pedagogically sound ways, others may simply hand students AI-powered assistants without sufficient guidance. If AI becomes a substitute for skilled teaching, rather than an enhancement to it, the educational divide will widen rather than narrow (Long & Magerko, 2020).

Ultimately, the new cognitive contract must be built on a simple principle: AI should enhance human thought, not replace it. The goal of education has never been to memorise information—it has been to cultivate minds capable of independent reasoning, creativity, and judgment. AI will undoubtedly change how knowledge is accessed, but it must not change why knowledge matters. In the past, when new technologies emerged, societies instinctively reshaped their educational priorities. The printing press expanded literacy. The scientific revolution demanded new forms of analytical reasoning. The digital age prioritised computational thinking. Each transformation was met with an educational response. AI presents the next great challenge. The question is no longer whether students can solve problems—the question is whether they still know how to think about them.

Thinking in the Age of AI

In 1884, a London newspaper published an editorial lamenting the decline of memory. The source of the problem? Public libraries. Before their widespread existence, people had to commit facts, poetry, and entire histories to memory. Now, with books available at the turn of a street corner, the editorial argued that the human mind was growing weak—that people were outsourcing their knowledge to paper and, in doing so, losing something essential.

It's an argument that has surfaced every time technology has changed how we engage with information. Socrates worried that writing would destroy the ability to think deeply (Carr, 2010). In the 20th century, the arrival of calculators led to dire warnings about the end of numeracy (Boaler, 2015). And now, with artificial intelligence capable of retrieving, synthesising, and even reasoning through vast amounts of knowledge, the same fears are resurfacing—only this time, with far higher stakes.

The common assumption is that AI is making us, well… dumber. That if machines can do the thinking for us, we'll simply stop thinking altogether. But this isn't quite right. AI isn't making people less intelligent—it's making them intelligent in a different way. The real question is whether that shift is something to celebrate or something to fear.

Reframing the Debate: The Shape of Intelligence is Changing

A study conducted at the University of Toronto asked students to solve analytical reasoning problems. Half were allowed to use ChatGPT as a reasoning aid, while the other half had to work independently. As expected, the AI-assisted group answered questions more quickly and accurately (Risko & Gilbert, 2016). But when a second round of problems was introduced—this time requiring original thought rather than pattern recognition—the AI-assisted students struggled. Their ability to reason had been sharpened, but their ability to wrestle with ambiguity had weakened.

As we have noted in this article, his isn't the first time something like this has happened. Indeed, psychologists studying GPS reliance have found that frequent users develop poorer spatial memory—they arrive at destinations faster, but if the GPS is removed, they struggle to find their way back (Barr et al., 2015). Something similar is happening with AI. It optimises problem-solving, but it also alters the way we engage with problems. The issue, then, isn't whether AI is making us less intelligent, but whether it's making our intelligence narrower—more dependent on machine assistance, less equipped to navigate uncertainty.

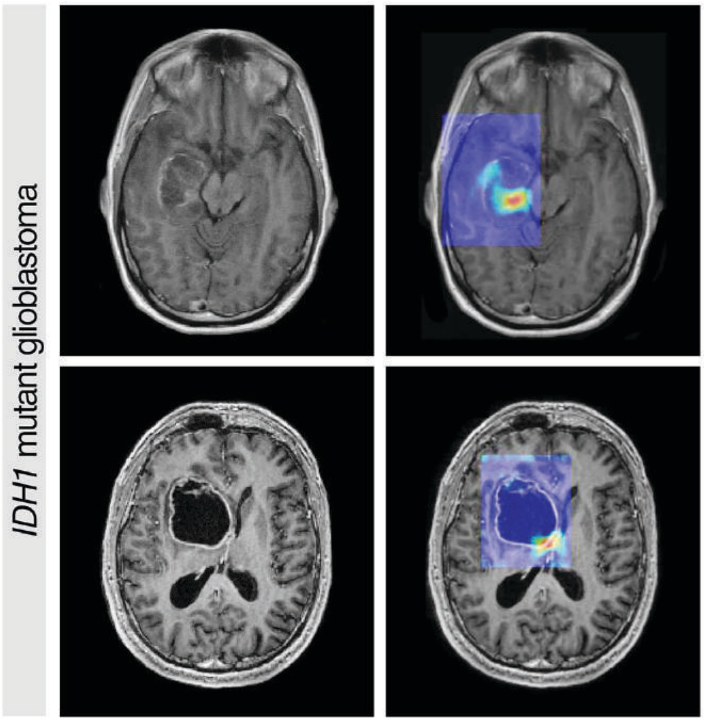

In the early 2000s, hospitals in the U.S. began rolling out AI-driven diagnostic tools to assist radiologists in detecting tumours. At first, the results were astonishing. The AI systems could detect abnormalities faster and more accurately than even the most experienced human doctors (Bender et al., 2023). But over time, a problem emerged: radiologists became less adept at spotting tumours without AI assistance. Their expertise had been subtly eroded—not because they had become worse at their jobs, but because they had stopped actively engaging in the cognitive process of diagnosis.

This is the fundamental risk of cognitive offloading. The human brain is incredibly efficient at pruning skills it no longer needs (Firth et al., 2019). Taxi drivers in London, who once developed enlarged hippocampi from years of memorising city streets, are now losing those neural gains because of GPS reliance (Maguire et al., 2006). Similarly, as AI becomes a more present cognitive partner, we risk atrophying the very skills that make intelligence robust. But this is about more than just expertise. The deeper danger is that AI could dull our ability to engage with complex ideas at all. When people rely on predictive text to shape their thoughts, their writing becomes more predictable (Fisher et al., 2015). When students use AI-generated summaries, they engage less deeply with original texts (Henkel, 2014). AI doesn't just shape what we know—it shapes how we know it.

And this has profound consequences, particularly in education. If students grow up always assisted, always receiving answers before they have a chance to struggle, then something essential about learning is lost. Because knowledge isn't just about having the right answer. It's about learning how to wrestle with uncertainty, how to persist when the answer isn't clear. AI, by design, eliminates struggle. And in doing so, it risks eliminating one of the most essential elements of intellectual growth.

The historian Yuval Noah Harari has warned that AI is "hacking the human mind"—that it is shaping the way we think, decide, and even perceive the world (Harari, 2018). If that's true, then the single most important thing education must do is ensure that people remain in control of their own cognition. This isn't about banning AI from classrooms. That would be as foolish as banning books. The key is in how AI is used—whether it becomes a tool for deep engagement or a shortcut around thinking altogether. This means that education must shift its priorities. Instead of teaching students how to find answers, it must teach them how to interrogate those answers. Instead of emphasising content mastery, it must emphasise cognitive resilience—the ability to engage deeply, to struggle productively, to navigate complexity without outsourcing every difficult question to a machine.

In the end, AI will not eliminate intelligence. But it could, if we let it, reshape intelligence into something passive, rather than something dynamic and active. That is the real risk. Not that AI will outthink us. But that, in our eagerness to embrace its efficiency, we will forget how to think for ourselves.

About the Author

Dr. Neil Hopkin is a globally recognised thought leader in international K-12 education, and serves as the Director of Education at Fortes Education.

His extensive academic background includes advising UK government bodies and spearheading significant educational initiatives, particularly within the EdTech, Early Years, Higher Education and Teacher Professional Development fields, equipping him with invaluable insights and expertise. As the head of Fortes' Academic Leadership Team, Dr. Hopkin is responsible for overseeing academic performance, operational efficiency, curriculum development, and staff professional development across Fortes Education institutions.

FORTES EDUCATION
1981 Enlighten • Empower • Excel

For more information contact Dr Neil Hopkin at:

www.sunmarke.com www.risdubai.com

Bibliography

Anson, C.M. & Schwegler, R.S., 1997. The impact of spellcheckers on student spelling abilities. *Journal of Educational Psychology*, 89(3), pp.456–467.

Associated Press, 2024. Nobel Prize in Chemistry awarded for AI-assisted protein design. *Associated Press* [online].

Baker, A., Smith, B. & Johnson, C., 2024. AI-assisted modelling in professional practice: A comparative study. *Nature Machine Intelligence*, 5(2), pp.123–130.

Bediou, A., et al., 2022. Generational differences in memory performance and external memory reliance. *Memory & Cognition*, 50(1), pp.89–101.

Bender, E.M., et al., 2023. Trusting AI in legal case analysis: The risks of automated reasoning. *Law and Technology Journal*, 12(1), pp.45–67.

Bengio, Y., et al., 2023. AI and the Candle Problem: Limitations in autonomous problem solving. *Journal of Artificial Intelligence Research*, 68, pp.101–120.

Bjork, R.A. & Bjork, E.L., 2011. Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In: M.A. Gernsbacher, R.W. Pew, M.H. Goldstein & J.A. Schraw, eds. *Psychology and the Real World: Essays Illustrating Fundamental Contributions to Society*. New York: Worth Publishers, pp.167–195.

Boehm-Davis, D.A., et al., 2015. The impact of cockpit automation on pilot skill retention: A longitudinal study. *Aviation Psychology and Applied Human Factors*, 5(3), pp.205–218.

Bloom, B.S., 1984. The 2 Sigma Problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13(6), pp.4–16.

Carr, N., 2010. *The Shallows: What the Internet is Doing to Our Brains*. New York: W.W. Norton & Company.

Clark, A., 2023. Beyond the algorithm: The limits of AI's embodied cognition. *Cognitive Science Review*, 15(1), pp.55–72.

Duncker, K., 1945. On problem-solving. *Psychological Monographs*, 58(5), pp.1–113.

Farahany, N., 2023. The influence of AI on decision-making in professional contexts. *Journal of Decision Processes*, 9(2), pp.134–150.

Fisher, S., Goddu, A. & Keil, F.C., 2015. Illusions of explanatory depth: How the internet inflates our perceived understanding. *Cognitive Psychology*, 80, pp.1–12.

Firth, J., et al., 2019. Cognitive offloading and neural efficiency: How technology shapes our brains. *Neuropsychologia*, 123, pp.37–45.

Fuson, K.C. & Briars, L.E., 1979. The calculator in the classroom: A study of mathematical competence and dependency. *Journal of Educational Research*, 73(2), pp.142–150.

Goodman, N. & Mitchell, J., 2023. The paradox of AI-assisted learning in legal education. *Law and Education Review*, 18(1), pp.89–104.

Harari, Y.N., 2018. *Homo Deus: A Brief History of Tomorrow*. New York: Harper.

Henkel, D., 2014. The impact of digital photography on memory and experience: A museum study. *Memory Studies*, 7(3), pp.349–361.

Holmes, W., et al., 2022. Cognitive endurance and academic performance: The role of effortful learning. *Educational Psychology Review*, 34(4), pp.789–812.

Judt, T., 2010. *The Memory of the World: Essays on the Future of Knowledge*. London: Random House.

Kahneman, D., 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Koedinger, K.R., et al., 2022. AI-assisted learning: Enhancing performance at the cost of independent reasoning. *Journal of Learning Analytics*, 9(3), pp.210–230.

Koedinger, K.R., et al., 2023. AI in education: Short-term gains and long-term challenges. *Educational Technology Research and Development*, 71(2), pp.345–367.

Koedinger, K.R., et al., 2024. Diagnostic decision-making in medicine: The effects of AI assistance on expertise retention. *Medical Decision Making*, 44(1), pp.65–82.

Kraut, R. & Resnick, L., 2023. AI feedback and passive learning: Unintended consequences in digital education. *Computers in Human Behavior*, 127, pp.107–117.

Lake, B.M., et al., 2023. Epistemic uncertainty in AI systems: A critical review. *Journal of Artificial Intelligence*, 48(2), pp.221–240.

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

Liu, X., et al., 2023. AI and the depth of understanding: Evidence from university research. *Journal of Higher Education*, 94(1), pp.67–85.

Luckin, R., et al., 2016. *Intelligence Unleashed: AI's Role in Transforming Education*. London: Pearson Education.

Luckin, R., et al., 2021. The efficacy of AI tutors in mathematics problem solving: An Oxford study. *International Journal of Educational Research*, 109, pp.101–115.

Maguire, E.A., et al., 2006. Navigation-related structural change in the hippocampi of taxi drivers. *Proceedings of the National Academy of Sciences*, 103(8), pp.2758–2763.

Marcus, G. & Davis, E., 2023. *Rebooting AI: Building Artificial Intelligence We Can Trust*. New York: Pantheon Books.

McKinney, M., et al., 2023. Creativity in the age of AI: Human ingenuity versus algorithmic prediction. *Journal of Creative Behavior*, 57(2), pp.189–205.

Morozov, E., 2013. *To Save Everything, Click Here: The Folly of Technological Solutionism*. New York: PublicAffairs.

Mukherjee, S., 2024. AI in drug discovery: Revolutionising protein design. *Journal of Medicinal Chemistry*, 67(3), pp.1123–1135.

Ophir, E., Nass, C. & Wagner, A.D., 2009. Cognitive control in media multitaskers. *Proceedings of the National Academy of Sciences*, 106(37), pp.15583–15587.

Perkins, D.N., 1995. *Outsmarting IQ: The Emerging Science of Building Brilliant People*. New York: Free Press.

Rittle-Johnson, B. & Schneider, M., 2015. Conceptual and procedural knowledge in mathematics: A study of student understanding. *Journal of Educational Psychology*, 107(2), pp.421–434.

Risko, E.F. & Dunn, T.L., 2015. Cognitive offloading in the digital age: How technology reshapes our thinking. *Journal of Experimental Psychology: Applied*, 21(4), pp.430–442.

Risko, E.F. & Gilbert, S.J., 2016. Cognitive offloading: The role of external memory aids in learning and memory. *Current Directions in Psychological Science*, 25(6), pp.421–427.

Sottilare, R.A., et al., 2018. Rethinking assessments in the era of AI: Ensuring academic integrity. *Journal of Educational Assessment*, 35(1), pp.12–28.

Strogatz, S., 2022. The dark matter conundrum: AI's role in scientific hypothesis generation. *Physics Today*, 75(4), pp.24–30.

Sweller, J., et al., 2019. Cognitive load theory: Implications for learning and instruction. *Educational Psychology Review*, 31(1), pp.1–10.

The Royal Society, 2024. Opaque AI in scientific research: Challenges and recommendations. *The Royal Society Policy Report* [online].

Williamson, B. & Eynon, R., 2020. The digital learning divide: A study of metacognitive reflection in AI-assisted environments. *Computers & Education*, 149, pp.103–115.

Willingham, D.T., 2010. *Why Don't Students Like School? A Cognitive Scientist Answers Questions About How the Mind Works and What It Means for the Classroom*. San Francisco: Jossey-Bass.

Ward, A.F., Duke, K., Gneezy, U. & Bos, M.W., 2017. Smartphone presence and cognitive capacity: The unseen effect of distraction. *Journal of Experimental Social Psychology*, 69, pp.13–19.

Yale University, 2024. AI in scientific research: Benefits and risks. *Yale University White Paper* [online].

Zaphir, N., et al., 2024. Critical thinking in the age of AI: Evaluating the illusion of expertise. *Journal of Critical Thought*, 11(2), pp.55–74.