


Disappearing Thought

**What MIT's Study Reveals
About AI and Children**



Dr Neil Hopkin
Director of Education
Fortes Education

Introduction: The Neural Price of Convenience

In the spring of 2024, researchers at the Massachusetts Institute of Technology published an unsettling preprint that drew little initial fanfare beyond the corridors of cognitive neuroscience. Their study appeared deceptively simple: wire up students with EEG headsets, set them the familiar task of writing essays, and compare their brain activity when they worked unaided versus when they leaned on ChatGPT. The results were stark. In the AI-assisted condition, neural connectivity in regions associated with memory and creativity dropped precipitously. Worse still, when those same students later tried to recall what they had written, they faltered; many could not paraphrase their own essays, let alone reconstruct key arguments (MIT Media Lab, 2024). In a single stroke, the study suggested that artificial intelligence might not only lighten the cognitive load of writing, it might hollow it out.

This revelation unsettles educators and neuroscientists alike because it touches on something both fragile and formative: the architecture of developing brains. While the adult professional might view ChatGPT as an efficiency tool, a way of pruning the tedium of emails or generating draft frameworks for reports, the child or adolescent is in a different cognitive landscape altogether. Their neural connections are not yet fixed highways but malleable trails, carved through repeated struggle, retrieval, and practice. If those trails are bypassed too often, do they ever fully form?

It is precisely this question that frames a growing conversation among thinkers across neuroscience, psychology, ethics, and education. Terry Sejnowski, a pioneer of computational neuroscience, has long argued that the gold standard for learning remains one-to-one tutoring. When asked about AI's potential, he did not recoil from the technology but pointed out its opportunity: if used correctly, large language models could approximate the personalised attention of a tutor, stimulating curiosity, posing questions, and drawing learners into dialogue. The danger, in his view, is not the presence of AI but its misuse, as a substitute for thought rather than a provocation towards it (Sejnowski, 2023).

By contrast, Pattie Maes, one of the MIT study's senior authors, has sounded a sharper alarm. Her concern is not just about learning outcomes in the narrow sense of exam performance, but about what she calls "the natural intelligence of society" (Maes, 2024). If the neural scaffolding of memory and creativity weakens through habitual outsourcing, the implications stretch well beyond the classroom. They touch on our collective capacity to sustain democracy, innovation, and even cultural memory. Maes' intervention serves as a reminder that the classroom is not a silo but a node in a much larger social brain.

Henry Shevlin, a philosopher at the University of Cambridge, counsels against premature panic. He is quick to point out that the MIT work is a preprint, with a modest sample, and that the story of educational technology is littered with moral panics that later faded. What matters, he argues, is how AI is woven into pedagogical practice. Used judiciously, it may enhance engagement and creativity; used carelessly, it may corrode them (Shevlin, 2024). Shevlin's nuance is important, not least because it forces us to look past headlines towards a subtler truth: AI is neither panacea nor poison, but a tool whose effects depend on the context of use.

Other scholars add further complications. Aditya Vashistha and Mor Naaman, in their work at Cornell, have shown how AI's autocomplete and suggestion features can homogenise writing and thought, nudging individuals towards convergent outputs that feel original but are, in fact, the product of hidden steering (Vashistha & Naaman, 2023). Their research dovetails uneasily with the MIT findings: even if neural connectivity were unaffected, the quiet narrowing of imagination might still represent a profound educational loss. And Susan Schneider, a philosopher of mind and ethicist of AI, takes the issue into childhood directly. For her, the core concern is developmental ethics: not whether adults should be free to use AI to shave minutes off their working day, but whether children, whose executive function and critical faculties are still plastic, should be exposed without guardrails (Schneider, 2024).

Taken together, these voices form less a consensus than a conversation, a roundtable in which neuroscientists, philosophers, and technologists circle the same problem from different angles. The MIT study is the spark, but the fuel comes from wider debates that have been smouldering

for over a decade: Nicholas Carr's anxiety that the internet erodes deep reading (Carr, 2010), Maryanne Wolf's insistence that the "reading brain" is a cultural achievement constantly at risk of reversal (Wolf, 2018), Alison Gopnik's reminders that play and exploration are not luxuries but evolutionary necessities in childhood (Gopnik, 2020). What AI has done is collapse these concerns into a single, measurable phenomenon: the possibility that children's neural engagement can be visibly and quantifiably reduced by the presence of a generative model.

The challenge for education, then, is not merely technical but philosophical. Should schools position AI as a tutor, a tool, or a threat? How do we prevent the "cognitive offloading" measured in the MIT lab from becoming a cultural default? Can we design practices that keep the child's brain in the driver's seat, with AI as a navigator rather than a chauffeur? These questions matter because they go to the heart of what it means to learn in an age when thinking itself can be outsourced.

This article seeks to draw these threads together in a sustained inquiry. It is not at all a manifesto against AI, but neither is it a cheerleading exercise for its adoption. Instead, it will trace the contours of the debate by juxtaposing and weaving the voices of primary commentators such as Sejnowski, Maes, Shevlin, Vashistha, Naaman, and Schneider with the broader scholarship of neuroscience, developmental psychology, pedagogy, and educational policy. Thirty or more authors will anchor the analysis, from Stanislas Dehaene's studies on neural coding in literacy, to John Sweller's work on cognitive load, to Sonia Livingstone's research on digital childhoods.

By structuring this exploration as a conversation, rather than a monologue, my aim is to capture both the urgency and the complexity of the issue. If the MIT study revealed anything, it is that the question of AI and childhood is not simply about performance metrics but about the conditions under which neural connections form, strengthen, and endure. It is a question of architecture—how the scaffolding of thought is built, and whether AI is to be the craftsman's tool or the wrecking ball.

The discussions that follow will move from the neuroscience of cognitive load, through the ethics of developmental vulnerability, to the global policy landscape that shapes how schools respond. Along the way, we will return repeatedly to the core tension: how to harness the generative power of AI without sacrificing the generative power of the child's brain.

Cognitive Load and Neural Connectivity

The quiet recorded in the MIT Media Lab was not the calm of rest, but the eerie silence of underuse. Students who wrote with the aid of ChatGPT produced fluent essays, yet their brains, measured through EEG, showed a collapse of connectivity in the regions linked to creativity and memory. When asked later to recall what they had written, many faltered. The text existed, but it had not passed through the deeper circuits of thought and recollection. For teachers, the result feels strangely familiar: the immaculate homework that seems untouched by struggle, the polished answer that carries no trace of the child's voice. What the electrodes captured in real time was a phenomenon every educator has seen: learning without labour, output without imprint.

The paradox has haunted pedagogy for decades: how much help is too much? John Sweller's Cognitive Load Theory offers a starting point. Working memory, he argued, is a fragile bottleneck; overload it with extraneous detail and students drown. But the shadow side of this truth is rarely acknowledged. Strip away too much, and there is no work left to be done. Sweller himself acknowledged that reducing load is not synonymous with increasing learning (Sweller, 1988). The MIT study seemed to catch this corollary in the act. The students' brains had not been taxed, and without cognitive strain there was nothing left to consolidate.

Stanislas Dehaene's neuroscience deepens the picture. His account of "neuronal recycling" describes how the brain re-purposes ancient circuits for modern cultural inventions like reading (Dehaene, 2009). That repurposing is effortful. Children learning to decode words engage their visual systems in painstaking repetition, carving new pathways across cortical regions. To remove the labour is to short-circuit the process. Neural connections are not delivered wholesale; they are

sculpted through failure, rehearsal, and eventual mastery. The muted EEG signals in MIT's study may be the physiological echo of a skipped rehearsal. What looks like efficiency may in fact be impoverishment.

Maryanne Wolf's argument for the "deep reading brain" offers a cultural parallel to Dehaene's biology. In *Reader, Come Home*, she warns that digital skimming erodes the circuits responsible for empathy, reflection, and critical analysis (Wolf, 2018). These higher-order capacities emerge only through the sustained cognitive labour of deep reading. Wolf's fear is not about comprehension in the narrow sense but about the broader loss of the mental habits that anchor a literate society. When her warning is placed beside the MIT findings, the resonance is clear. Whether through skimming or outsourcing, the danger is the same: cognition stripped of its effortful depth. Dehaene, Sweller, and Wolf, from their different disciplines, converge on a single principle: neural architecture is laid down in the struggle.

Yet struggle alone is not enough. Terry Sejnowski, whose career has bridged neuroscience and machine learning, insists that the miracle of one-to-one tutoring lies not in leaving students to flail, but in carefully calibrated guidance. Tutoring reduces wasted effort while preserving the essential challenge (Sejnowski, 2023). In this sense, AI holds promise. If designed to probe, question, and nudge, it could mimic the conditions under which brains wire most effectively. But this is not how AI is typically encountered in classrooms. Too often, it functions as an answer engine, offering fluency without friction. The MIT results suggest what happens when the balance tips towards removal of effort rather than its refinement.

The motivational dimension makes the paradox sharper still. Mary Helen Immordino-Yang has demonstrated that cognition and emotion are deeply entwined, with affective engagement fuelling the consolidation of neural pathways (Immordino-Yang, 2016). When learners care, the connections endure. If AI strips tasks of their personal investment, by doing the "interesting" work on their behalf, then the collapse in brain activity is as much emotional as cognitive. Michael Posner's research on attention dovetails here. Attention is not simply the gatekeeper of cognition; it is a trainable network that develops gradually in childhood (Posner & Rothbart, 2007). If children are invited too often to look away because the machine is looking for them, then the attentional scaffolding itself weakens. What MIT measured may therefore be a compound deficit: less thinking, less caring, less attending.

Consider, then, what happens when we turn to the role of retrieval. Henry Roediger and Jeffrey Karpicke famously showed that memory does not strengthen through exposure alone, but through the act of recall (Roediger & Karpicke, 2006). Struggle to retrieve, and the memory trace consolidates. If AI becomes the recall system, summoning facts, structuring arguments, even generating counterpoints, the student may miss the very exercise that cements learning. This is precisely why the MIT participants, when asked to recall their own AI-assisted essays, found themselves adrift. Without the act of retrieval, memory has no reason to fire, and without firing, it has no chance to wire.

Henry Shevlin urges us not to rush to apocalypse. As he points out, the MIT paper is a preprint with a modest sample, and educational technology has weathered moral panics before (Shevlin, 2024). His caution is necessary, not least because neuroscience can seduce with its imagery of glowing brains and collapsing networks. Yet Shevlin's reminder of context does not erase the broader worry. The danger is less about the exact magnitude of the drop than about the trajectory it suggests. If the default use of AI is to replace rather than provoke thought, the risks accumulate silently, one skipped retrieval at a time.

Daniel Willingham's aphorism distills the whole argument into a phrase: memory is the residue of thought (Willingham, 2009). The MIT study was a case study in residue that never formed. Essays flowed, but thought did not settle into memory. And here lies the uncomfortable truth. The problem is not that AI makes tasks too easy in a general sense, but that it makes them easy in precisely the domain where difficulty is the point. Writing, reading, and reasoning are not chores to be eliminated; they are the exercises through which the brain constructs itself. To remove the labour is to remove the growth.

What emerges across these perspectives is not a rejection of AI but a recalibration of its role. Sweller, Dehaene, Wolf insist on the necessity of struggle. Sejnowski reframes that struggle as something to be carefully guided, not erased. Immordino-Yang and Posner reveal how motivation and attention are inseparable from cognition. Roediger and Karpicke remind us that memory is built in the act of retrieval. Shevlin tempers alarm but leaves the central paradox intact. And Willingham offers the line that crystallises the whole: no thought, no memory.

The electrodes in the MIT lab did not merely capture the absence of brain activity; they captured the cost of misplaced convenience. The task ahead is not to exclude AI but to choreograph it so that the struggle remains intact. Because the neural architecture of childhood is not built on answers, but on the effort to find them.

AI as Tutor vs. AI as Shortcut

The question, after the MIT study, is not whether AI diminishes neural activity but whether it has to. The silence of the electrodes revealed one possible future, a future in which convenience drains the mind of effort. But there is another possibility: that AI could become not a substitute for thought but its provocation. The distinction is subtle but decisive. Is AI the chauffeur, carrying the student smoothly to the destination, or is it the tutor walking beside them, insisting that they do the work of every step?

Terry Sejnowski has long been clear on this point. The value of one-to-one tutoring lies not in removing difficulty but in calibrating it, keeping learners suspended in what Vygotsky once called the “zone of proximal development” (Vygotsky, 1978). The tutor does not provide the answer; they nudge, question, and challenge until the learner’s own circuits ignite. Sejnowski suggests that large language models could, at their best, approximate this dance. They could be Socratic companions, pressing the learner to explain, retrieve, and reflect (Sejnowski, 2023). If that vision were realised, the very problem caught in MIT’s EEG traces, the collapse of brain connectivity, might be reversed.

Benjamin Bloom would have recognised this possibility. His “2 sigma” problem, first articulated in 1984, showed that one-to-one tutoring could lift student performance by two standard deviations above conventional classroom teaching (Bloom, 1984). The holy grail of education since then has been to replicate that effect at scale. For decades, the answer proved elusive; tutoring was too resource-intensive. Now, with AI, the question is back on the table. If machines can replicate not just the answers but the probing rhythm of a tutor, they might finally offer Bloom’s two sigma to every child.

But here lies the danger. Most current uses of AI in classrooms do not resemble Bloom’s ideal or Sejnowski’s vision. They resemble a shortcut. A child asked to write a paragraph may type “Explain photosynthesis at Year 7 level” and watch as the machine generates a polished answer. This is not tutoring. It is task completion. Henry Shevlin’s caution about overreacting to the MIT study is well taken (Shevlin, 2024). Yet his own distinction underscores the risk. What matters is not whether AI is present, but how it is woven into pedagogy. At present, too often, it is woven as substitution, not stimulation.

This distinction becomes more urgent when we consider the ethics of childhood. Susan Schneider, writing from the perspective of philosophy of mind, warns that children are uniquely vulnerable to cognitive offloading (Schneider, 2024). Unlike adults, whose neural scaffolding is largely in place, children are still laying down the very architecture of their executive functions. To displace that effort with AI is to weaken the circuits before they are built. Schneider is not advocating prohibition, but boundaries: age-appropriate exposure, deliberate design, and constant reflection on the developmental stakes.

The pedagogy of scaffolding provides a parallel here. Linda Darling-Hammond has long argued that effective teaching is about sequencing challenge, offering support until mastery is achieved, then withdrawing it to allow independence (Darling-Hammond, 2008). AI can, in principle, provide

this kind of scaffolding with exquisite responsiveness. But if it simply provides finished answers, it collapses the scaffold before the learner has climbed. Carol Dweck's work on mindset sharpens this worry (Dweck, 2006). If children learn to associate AI with effortless success, they risk cultivating what she would call a fixed mindset: an expectation that competence arrives without struggle. A tutoring AI would encourage persistence, error, and growth. A shortcut AI would train dependence.

This is not an abstract distinction; it can be seen in how children interact with feedback. Dylan Wiliam has shown that feedback is only effective when it activates the learner's own thinking (Wiliam, 2011). AI that simply marks work or supplies corrections is feedback of the least useful kind: information that bypasses thought. By contrast, AI that asks the learner to reflect on their reasoning, to attempt again, to justify a choice, aligns with Wiliam's principle. The difference is not in the presence of AI but in whether it serves as a mirror or a crutch.

At stake here is the nature of agency. Jerome Bruner once described learning as the process of "going beyond the information given" (Bruner, 1973). AI can either suppress or stimulate this going beyond. In its shortcut form, it gives students the information and invites them to stop. In its tutoring form, it withholds just enough to invite the leap. The challenge for educators is to choreograph AI's role so that it leans toward provocation. Otherwise, the seductive efficiency of shortcut use may harden into habit, eroding the very circuits that Bloom, Bruner, and Sejnowski saw as the engines of learning.

The paradox becomes sharper still when we consider assessment. If AI is used to complete essays, then assessment ceases to measure thought; it measures prompt engineering. In such a world, retrieval, the act that Roediger and Karpicke showed to be foundational to memory (Roediger & Karpicke, 2006), is outsourced. The shortcut does not just save time; it displaces the act of remembering. Here again, Willingham's dictum returns like a refrain: memory is the residue of thought (Willingham, 2009). Without thought, there is no residue. Without residue, the work of school dissolves.

Yet one must resist fatalism. If shortcuts are one trajectory, tutoring is another. The choice is not made by the technology but by its design and deployment. In pilot classrooms, teachers who instruct students to draft first, then use AI for critique, find higher engagement and stronger recall. The machine in this configuration is not the author but the interrogator. It is the difference between a pupil handing in an answer sheet filled by someone else, and a pupil wrestling with their own draft under a tutor's questioning gaze. One path leads to silence on the EEG, the other to sparks of connectivity.

The debate, then, is not between those who embrace AI and those who resist it. It is between two visions of what AI is for. Sejnowski, Bloom, Darling-Hammond, Dweck, Wiliam, Bruner—all converge, implicitly, on the same vision: AI as provocation. The shortcut vision is seductive, but it is also sterile. The tutoring vision is demanding, but it holds the possibility of preserving, and perhaps even amplifying, the neural struggle that makes learning real.

The MIT study showed us what happens when AI is allowed to drift into the role of shortcut. The challenge now is to design, regulate, and teach towards the alternative. For the stakes are not only about grades or essays, but about whether the next generation learns to think with machines or only through them.

Homogenisation and Hidden Cognitive Risks

The first time you watch a child produce an essay with AI, it is hard not to be impressed. Where there might once have been hesitations, crossings-out, and awkward phrasing, there is now polish. Sentences arrive smoothly, the grammar flawless, the structure clear. Teachers who have battled for years with half-formed paragraphs may even feel a quiet relief: at last, something coherent, something easy to mark. Yet beneath the satisfaction lurks a quieter danger. Fluency is not the same as originality. And in learning, polish can be the enemy of growth.

Researchers such as Aditya Vashistha and Mor Naaman have already demonstrated what intuition tells us: predictive systems funnel students towards a narrow band of expression. Ideas converge, arguments flatten, the quirky edges of human voice are sanded down. But to think of this merely as a stylistic matter would be to miss the deeper risk. Nicholas Carr, more than a decade ago, argued that the internet did not just distract us; it rewired us, training our minds to skim, to prefer speed over depth. What the predictive model introduces is a new twist: not just speed, but convergence. The prose looks correct precisely because it is statistically the most likely thing to say. When children lean too heavily on these models, their writing ceases to be an act of invention and becomes instead a performance of probability.

Maryanne Wolf's concern for the "deep reading brain" provides another angle into this paradox. Her work shows that empathy, inference, and critical thought are cultivated only through the slow, effortful circuits of sustained reading. Strip away the slowness, and the deeper capacities wither. In a similar way, writing that is made too easy loses its generative power. The act of struggling for a metaphor, of reshaping a sentence until it fits, is not wasted effort. It is the very process through which thought deepens. When AI supplies the phrasing, the struggle vanishes, and with it the growth. Sherry Turkle captures the cultural consequences of this shift. She has shown how children who converse with robots begin to internalise those interactions, expecting the same smoothness and predictability from people. What happens when children compose with AI is parallel: the uniform cadences of the machine become their own. The surprise of authentic voice, the awkward turn of phrase that signals genuine thought, is replaced by the polished cadence of the model. In time, even the expectation of eccentricity fades. Classrooms fill with fluent but interchangeable essays, and no one notices how much has been lost.

Neil Postman warned long ago that every medium carries its own epistemology. Television reshaped public life into entertainment, he argued; the danger of predictive text is that it reshapes it into convergence. The logic of the medium rewards what is fluent and penalises what is difficult. Cathy O'Neil shows us how such logics are rarely neutral. Algorithms, for all their apparent objectivity, encode bias and replicate inequities. If children learn to trust their essays to AI, then their voices are being shaped by systems whose parameters they cannot see, and whose blind spots they cannot challenge. Shoshana Zuboff extends this worry into the domain of power: children's words themselves become data, harvested, commodified, and fed back into systems optimised for engagement and control. The danger of homogenisation is thus not only cognitive but cultural and commercial. And yet the paradox is that the same predictive capacities that narrow could also, if reconfigured, expand. There is no necessity that a model always supply the most probable answer. An AI trained to offer eccentric alternatives, to suggest metaphors from unexpected domains, to withhold the obvious phrasing and insist on elaboration—such a tool could provoke, not suppress, originality. The problem is not intrinsic to the technology but to its defaults, and to the ease with which we mistake fluency for learning.

This is where pedagogy becomes crucial. If teachers simply accept AI outputs as evidence of competence, homogenisation will become invisible, creeping across classrooms like ivy. Essays will be easier to grade, literacy scores may even climb, and all the while children's authentic voices will quietly thin. But if AI is deliberately positioned as a provocation—as a mirror against which students test their own drafts, or as a generator of wild analogies to be refined by human judgment—then the same system could widen expression. The danger is not that the machine speaks, but that we allow its speech to define what counts as good enough.

Homogenisation, then, is a subtler risk than neural silence. A silent brain announces itself; uniform thought slips under the radar. The student who hands in a polished AI-assisted essay may even be praised, the absence of originality hidden behind the gleam of fluency. But over time, the cost will become apparent. A generation of children may learn to write correctly without ever learning to write differently. The classroom, once a cacophony of voices, risks becoming a chorus, in tune but impoverished. To resist this, educators must resist the seduction of polish. Originality is noisy, awkward, sometimes even wrong. It is also the mark of a mind genuinely at work. The danger of AI is not only that it makes thinking too easy, but that it makes uniformity feel like success. If silence was the warning of Section I, sameness is the warning here. Both erode learning, one through absence, the other through conformity. And both require teachers to insist that children's work remain their own: effortful, eccentric, and gloriously resistant to prediction.

Developmental Ethics and Childhood Vulnerability

There is a comforting story adults tell themselves about technology: we can take the shortcuts and still be ourselves. We have earned it: years of drafts and dead ends, the muscle of recall built the slow way, so an autocomplete here or a summary there feels like changing the tool, not the hand. Children do not live inside that story. They are still assembling the hand: the habits, the attention, the patience, the will. Smooth the path too early and the path becomes them.

This is why the quiet trace in the MIT lab lands differently when you remember the age of the writers. A lull in neural engagement for a seasoned professional is a blip in the weather; for a ten-year-old it can be climate. Susan Schneider's insistence that childhood is not a smaller version of adulthood sits exactly here: executive functions, working memory, self-control, the very systems AI is so good at standing in for, are still under construction (Schneider, 2024). Take over the work and you are not merely helping; you are building the expectation that the bridge will always be there, so the brain never learns to span the gap.

Alison Gopnik's picture of the child as "little scientist" is the counterweight to that temptation. Childhood learning is not a clean route from question to answer but a wide field of hypotheses, false starts, and play, the evolutionary bet that exploration breeds flexible minds (Gopnik, 2009). A predictive model narrows that field by design. Its virtue is efficiency; its vice, for a learner, is the same. The wrong turn that would have become a memory gets bypassed. The odd analogy that would have stretched a concept never needs to be made. The classroom becomes a conveyor belt where it used to be a playground.

If that sounds sentimental, consider how memory is actually made. Mary Helen Immordino-Yang shows that emotion binds cognition; care and curiosity are not decoration but the glue that makes learning last (Immordino-Yang, 2016). Michael Posner's work on attention adds that the "spotlight" we bring to bear is itself trainable, a network strengthened through repeated demands (Posner & Rothbart, 2007). Now place those claims next to Willingham's plain sentence, memory is the residue of thought, and the developmental risk gains contour (Willingham, 2009). When AI does the interesting part, the child neither cares enough nor attends long enough to think hard enough. No care, no focus, no residue. The MIT electrodes did not just catch under-stimulation; they caught the slow unteaching of habits a child has not yet had time to build.

Jean Twenge's long arc on adolescence is less about screens as objects and more about displacement as a force (Twenge, 2017; 2024). Time that would have been spent in unsupervised play, friendship, sleep, or boredom, those awkward spaces where resilience grows, gets colonised by always-on systems with effortless rewards. Translate that logic into schoolwork and a second displacement appears: tasks once built for persistence become tasks built for completion. Jonathan Haidt's language of antifragility, whatever debates attend it, names the developmental hinge clearly: children need manageable stressors to become strong; insulation breeds fragility (Haidt, 2023). An AI that makes the work frictionless may feel like protection; in the long run it is closer to deconditioning.

But childhood is never generic, and the ethics cannot be either. Sonia Livingstone's research on digital childhoods reads as a quiet corrective to universal pronouncements (Livingstone, 2019). Context matters. A home that insists "think first, then ask the tool" turns the same system into a tutor; a home that celebrates fast polish trains dependence. Schools that have time and staff to hold oral checks, to ask for drafts before the machine is opened, will buffer risk; schools pressed for results may drift, under pressure and with the best of intentions, toward convenience. The developmental question is therefore also a distributive one: who gets the scaffolding that keeps AI as provocation, and who gets the shortcut that becomes habit?

Neuroscience does not let us off the hook. Stanislas Dehaene's account of neuronal recycling shows how literacy commandeers older visual pathways through effortful practice (Dehaene, 2009); Maryanne Wolf warns that the "deep reading brain", those long circuits for inference and empathy, is a cultural achievement, fragile enough to be reversed within a generation (Wolf, 2018). Neither claim is an argument against technology. Both are an argument for pacing: do not remove the very resistance through which those circuits take shape. If Section III worried about sameness

in the work children produce, the deeper worry here is sameness in the brains they are allowed to build.

Return for a moment to Schneider's point, now in conversation with Gopnik and Haidt. If play is the laboratory of childhood, and if small stressors are its curriculum, then the ethics of AI reduces to an unglamorous discipline: keep the hard part. Let the model check, not compose; quiz, not answer; suggest, not decide. This is not a list of classroom tricks but a stance about development. It is the difference between training wheels, supports designed to be removed, and a wheelchair, supports that quietly replace movement itself. The former strengthens the child's balance; the latter, misapplied, atrophies it.

The temptation, of course, is that the wheelchair looks like progress. A Year 5 class hands in crisp paragraphs; a parent sees fluent homework without the evening's argument; a dashboard records gains. Convenience is persuasive in the short run and expensive in the long run. You do not notice, term by term, that fewer pupils risk eccentric metaphors, or persist past confusion, or sit with a hard text without reaching for help. You notice later, when the older years need the stamina that was never trained because the tool was always there.

Livingstone's equity warning returns here with force. If the guardrails we are describing rely on adult time, teacher capacity, stable routines at home, and institutional slack, then the learners who most need friction will be the least likely to get it. That is not a reason to abandon AI; it is a reason to design expectations that make friction non-negotiable. Oral retrievals that cannot be outsourced. "Draft first, then tool" as a norm. Tasks that prize surprise over polish. These details belong properly to pedagogy, but their rationale is ethical: they protect the developmental goods that shortcuts erode.

It would be easy to turn this into a puritanical case for suffering. That is not what the research supports. Sweller's cognitive load still matters, too much difficulty and working memory buckles (Sweller, 1988). The craft is calibration: enough challenge to ignite the networks, enough guidance to keep the learner moving. This is where Sejnowski's tutoring vision is not a contradiction but a bridge. A machine that behaves like a good tutor, asking, probing, holding back, noticing when to press and when to pause, honours development rather than short-circuiting it (Sejnowski, 2023). The question is not "AI or no AI?" but "Which work remains the child's?"

If you strip this section down to a single sentence, it is this: growth needs friction, and childhood is the only time you can build the love of it. The electrodes in MIT's lab were a technical instrument; what they picked up was a human risk. Remove struggle and you remove pride. Remove attention and you remove memory. Remove play and you remove flexibility. The ethics of AI for children are therefore less about content filters and more about choreography. We are deciding, in quiet choices about prompts and policies, whether to raise a cohort adept at using tools or a cohort whose tools have quietly used them.

There is a way through. It asks for the patience adults reserve for things that do not pay off this week. It asks schools to defend awkwardness as a sign that something important is happening. It asks designers to build friction back in: defaults that quiz before they tell, interfaces that slow the quick fix, options that surface the odd instead of the obvious. It asks parents to tolerate less-polished homework in exchange for children who can speak from notes in their own words. None of this is dramatic. All of it is developmental ethics in practice.

And when you take the long view, the case becomes almost embarrassingly simple. Dehaene's circuits, Wolf's deep reading, Immordino-Yang's affect, Posner's attention, Gopnik's play, Haidt's antifragility, Livingstone's equity, Willingham's residue, Schneider's warning: they do not cancel each other; they complete each other. Together they describe a childhood in which technology is welcomed precisely to the point that it does not steal the work by which a mind is made.

Pedagogical Pathways Forward

The warnings matter little unless they can be translated into practice. Teachers do not need another lament about what children are losing; they need clarity on how to keep the essential difficulty intact. The paradox is simple to state and hard to enact: if the danger of AI is that it makes thinking too easy, then the pedagogy of the future must be the art of keeping the hard part.

This is where retrieval becomes more than an assessment tool. Henry Roediger and Jeffrey Karpicke demonstrated years ago that the act of recalling information strengthens memory far more than rereading it (Roediger & Karpicke, 2006). “Testing,” in their sense, is not about grading but about re-firing the neural circuits that consolidate knowledge. AI threatens this because it offers recall on demand. Why remember when you can regenerate? Yet AI could also be turned to the opposite end: a tutor that withholds the answer, insists on a first attempt, pushes a student to retrieve before it reveals. The retrieval effect does not vanish in the presence of AI; it vanishes only if teachers allow the machine to bypass it.

The same is true of metacognition. John Flavell introduced the term to capture how learners think about their own thinking (Flavell, 1979). The reflective pause, “What strategy did I use? Did it work? What will I do differently next time?”, is as crucial as the answer itself. A shortcut AI erodes this by providing answers too quickly. But an AI designed to press reflection, “Why did you choose this? What alternatives did you consider?”, could make metacognition more visible than ever. The paradox here is familiar by now: the same system that can stifle thinking can, in the right choreography, provoke deeper reflection than a human teacher with thirty pupils may have time for.

Feedback offers another case. David Carless has argued for “feedback literacy,” the capacity of learners to interpret and act upon comments rather than passively receive them (Carless, 2015). Dylan Wiliam’s work points in the same direction: feedback must cause thinking, not end it (Wiliam, 2011). AI can either collapse feedback into instant correction, red lines and fixed sentences, or it can expand it into dialogue. Imagine a model that does not correct the essay but asks: Which part of your argument feels strongest? Where is your evidence thinnest? The tool becomes not a red pen but a Socratic partner, shifting the responsibility for improvement back onto the learner.

And what of practice itself? Anders Ericsson’s theory of deliberate practice has shown that expertise does not emerge from repetition alone but from repeated, effortful refinement at the edge of competence (Ericsson et al., 1993). Carol Dweck’s mindset research echoes the same principle: growth requires a willingness to persist in the face of difficulty (Dweck, 2006). AI can undermine both if it normalises completion without struggle. But it can also be turned into a practice coach: generating new variations of a problem, stretching learners beyond the comfort zone, refusing to let them settle for fluency when mastery requires friction.

John Hattie’s synthesis of educational research, *Visible Learning*, makes the point bluntly: the most powerful effects on achievement are those that help students see and regulate their own learning (Hattie, 2009). AI is not on his list, but the implication is clear. The technology will matter only insofar as it is harnessed to make learning visible. Left unchecked, it makes learning invisible, answers appear but the process evaporates. Designed well, it can expose misconceptions, surface reasoning, show learners how they learn.

Paul Kirschner’s critique of minimally guided instruction is instructive here. He warned that discovery without scaffolding overloads working memory and leaves learners floundering (Kirschner et al., 2006). The mistake would be to see AI as a license for pure discovery, “go and ask the model”, without structure. The corrective is not to retreat to rote, but to choreograph guidance: AI that nudges, prompts, scaffolds, and then fades, just as Darling-Hammond and Bruner envisioned. Diana Laurillard’s conversational framework captures this rhythm: learning is a dialogue between teacher and learner, concept and practice (Laurillard, 2012). If AI is to play a role, it must take the part of the conversational partner, not the ghostwriter.

There are reasons for optimism here. Peter Brown, Henry Roediger, and Mark McDaniel, in *Make It Stick* (2014), show that learning strengthens through desirable difficulties—spacing, interleaving,

retrieval, and effortful practice. Each of these difficulties looks, on the surface, like inefficiency. Yet they are the soil in which long-term mastery grows. AI, misused, will strip them out: no need for spacing when you can generate on demand, no need for retrieval when you can ask again, no need for interleaving when you can autocomplete. The task of pedagogy is to put the difficulties back in, even when students resist. That resistance, after all, is the sign that learning is happening.

What ties all these strands together is Sejnowski's earlier reminder about tutoring. A good tutor does not remove the struggle; they orchestrate it. Bloom's "2 sigma" dream of tutoring at scale may yet be realised, but only if AI is forced to behave like a good tutor. That means resisting its default as shortcut and choreographing it as provocation. A model that offers problems rather than solutions, prompts rather than completions, variations rather than repetitions, this is the kind of tool that could strengthen rather than weaken neural connectivity.

None of this will happen by accident. Left to the market, AI will continue to optimise for speed, polish, and convenience. It will flatten originality, bypass retrieval, and seduce teachers into mistaking fluency for mastery. To turn it into pedagogy, educators must act with intention: design tasks where AI cannot be the author but can be the questioner; insist on drafts before polish; build routines where the tool's role is to provoke thinking, not replace it.

The future of pedagogy with AI is therefore not about whether schools permit or ban it. It is about whether schools defend the difficult. Retrieval, metacognition, feedback, deliberate practice, desirable difficulties, these are the "hard parts" that make learning real. If AI is to remain in classrooms, it must be harnessed to preserve them. Otherwise, the convenience that feels like progress will in fact be erosion.

Global and Policy Perspectives

The paradoxes that surface in the classroom are never sealed inside it. A child turning to AI for a draft is already responding to forces larger than their own choice. What teachers see as a struggle over pedagogy is, in truth, the echo of pressures set by systems, assessment frameworks, global agendas, policy incentives, that define what counts as learning. If AI is allowed to smooth away the difficult, it will not be only because the machine is capable of doing so, but because the culture of schooling makes that smoothing desirable.

When Pattie Maes reflected on the MIT findings, she did not stop at the individual brain. The question for her was what happens when under-stimulation scales up, when classrooms across countries embrace convenience, and the collective "natural intelligence of society" thins. At that point the debate ceases to be about technology in isolation and becomes a matter for policy. And here the tensions multiply.

Consider the global frameworks that promise preparation for the future. The OECD 2030 Learning Framework speaks of agency, adaptability, and creativity as essential dispositions. Andreas Schleicher has insisted repeatedly that nations must cultivate innovative learners who can navigate uncertainty. But the same accountability regimes that carry these ambitions measure success through predictable, standardised tests. The result is a quiet contradiction: schools are told to nurture creativity but are rewarded for conformity. In such an environment, the shortcut uses of AI—fast, polished answers, become a rational response, even when they corrode the very qualities policymakers say they want.

This is the tension Yong Zhao identifies when he warns that international competitiveness agendas often undermine genuine innovation. When every country chases the same benchmarks, diversity shrinks. Pasi Sahlberg extends the point to equity: technology is distributed unevenly, so reforms that assume equal access to opportunity risk deepening divides. Together, Zhao and Sahlberg cast Schleicher's optimism in a harsher light: the vision may be laudable, but the gravitational pull of systems drags practice in the opposite direction. The paradox of AI as tutor versus shortcut is mirrored by the paradox of policy that celebrates creativity but incentivises fluency.

UNESCO and the World Economic Forum add another register, speaking the language of ethics and human-centred design. AI in education must be transparent, equitable, trustworthy. Yet even here contradictions creep in. Transparency is urged, but few teachers or students can explain how a model generates its outputs. Creativity is praised, but the tools that spread most quickly are those that converge on predictable answers. “Human-centred” becomes a slogan rather than a standard when the pedagogy surrounding AI remains unchanged.

This cycle of optimism and contradiction is not new. Neil Selwyn has shown how every generation of educational technology, from television to tablets, was introduced with promises of transformation and ended by reinforcing existing patterns. Ben Williamson’s work reveals how ed-tech infrastructures embed logics invisibly, nudging teachers and students before anyone has debated what those nudges mean. Larry Cuban’s history of reform is full of such moments: grand claims, minor adaptations, the status quo intact. And Gert Biesta reminds us that the purpose of education cannot be reduced to outputs alone. Its task is also to cultivate subjectivity, judgment, and democratic agency, precisely the qualities most at risk if AI is allowed to settle into shortcut use.

Place these system-level reflections back alongside the neuroscientific and developmental warnings of Sejnowski and Schneider, and the full picture emerges. Neural silence, cognitive sameness, developmental fragility: none of these are inevitable, but all become more likely when systems reward ease over effort. Teachers may resist; individual classrooms may innovate. But policy pressure is relentless. If performance is measured in grades that AI can deliver faster, then convenience will trump pedagogy.

The alternative is not romantic rejection of technology but systemic defence of difficulty. Imagine assessment regimes that value originality more than polish, funding models that reward experimentation rather than standardisation, professional development that frames AI as questioner and challenger rather than as writer. Schleicher’s creativity, Zhao’s innovation, Sahlberg’s equity, Biesta’s subjectivity, when read together they amount to a single demand: that systems recalibrate what counts as learning so that AI can be harnessed for provocation rather than shortcut.

The electrodes in the MIT lab revealed how quickly brains can quieten when work is displaced. Policy frameworks show how quickly schools can drift when convenience is rewarded. Both point to the same conclusion: AI’s effect on education will not be settled by code alone. It will be determined by whether societies defend the slow, effortful, human work of thinking—or whether they reward its disappearance.

Conclusion – Between Silence and Sameness

What began in the MIT lab as a flicker on an EEG now stretches into a question for classrooms, systems, and societies. The students who wrote with ChatGPT produced fluent essays, but their brains fell quiet. Hours later, the words were gone from memory. It is tempting to treat this as an anomaly, a single data point in the noisy field of neuroscience. But as the arguments of scientists, psychologists, educators, and critics have accumulated, the pattern has taken shape. The danger of AI in education is not only that it makes thinking easier, but that it makes it too easy in precisely the domain where effort is the substance of growth.

The risks have shown themselves in two forms. The first is silence—the under-stimulation captured in the electrodes, the absence of retrieval, the erosion of circuits that require struggle to form. The second is sameness—the homogenisation Vashistha and Naaman identified, the polished convergence Nicholas Carr saw foreshadowed in the shallows of digital reading, the cultural narrowing that Turkle, Postman, O’Neil, and Zuboff each described from different vantage points. Silence starves the brain of activity; sameness starves it of originality. Both are erosions of learning, and both are easily mistaken for progress.

What has emerged across these reflections is a chorus of caution, not in the sense of prohibition but in the sense of care. Sweller’s cognitive load, Dehaene’s neuronal recycling, Wolf’s deep

reading, Immordino-Yang's emotion-cognition link, Posner's attention networks, Roediger and Karpicke's retrieval effect, Willingham's dictum about thought and memory, together they remind us that neural growth is built on friction. Sejnowski, Bloom, Darling-Hammond, Dweck, William, Bruner each insisted, in different ways, that learning requires orchestration of difficulty, not its removal. Schneider, Gopnik, Twenge, Livingstone, Haidt, each warned that childhood is fragile precisely because its habits are still forming, and that shortcuts in development are not neutral. Schleicher, Zhao, Sahlberg, UNESCO, Selwyn, Williamson, Cuban, Biesta, all pointed to the system level, where policy pressures risk rewarding the very behaviours that undermine thinking.

What ties these strands together is the paradox that the very features that make AI so attractive, its fluency, its efficiency, its predictive power, are the same features that threaten the architecture of learning if left unchecked. To give children fluency without struggle is to offer them food without digestion: it fills but does not nourish. To give them efficiency without originality is to invite them to walk on well-trodden paths but never to blaze their own.

And yet, woven through these warnings is another theme: possibility. Sejnowski's optimism about AI as a tutor, Bloom's 2-sigma challenge, Laurillard's conversational framework, Brown, Roediger, and McDaniel's "desirable difficulties," Ericsson's deliberate practice, all hint at a different trajectory. The same models that collapse thinking into silence or convergence could be re-designed to provoke, to stretch, to question, to withhold. The silence need not remain silence; the sameness need not remain sameness. The hinge is choreography: how schools, systems, and societies choose to frame the work that remains the child's.

This is why the ethical centre of the debate is not "AI or no AI" but "what do we let go of, and what do we insist stays hard?" The answer will not be the same for every subject or every age. A Year 2 child needs scaffolding a Year 12 student does not. Retrieval looks different in mathematics and in literature. Context matters: as Livingstone showed, children with strong scaffolding at home can risk more autonomy with AI; those without may need guardrails built in. Policy matters: as Zhao and Sahlberg warned, systems that prize testable outcomes will bend AI towards shortcuts, while systems that value originality and equity may hold open a space for provocation.

The temptation, always, will be to mistake polish for progress. Teachers exhausted by marking, parents desperate for homework completed without conflict, policymakers seeking metrics that rise, each has reason to welcome the fluency AI provides. But beneath that fluency, silence and sameness creep in. What must be defended is not simply content knowledge, but the conditions under which children learn to think, to care, to persist, and to create.

If there is a lesson to draw from this roundtable of voices, it is that education has always lived on the knife-edge of effort. Too much and the learner buckles; too little and they never grow. AI magnifies the stakes of that calibration. The electrodes at MIT gave us a glimpse of what happens when we slide too far towards ease. The uniform essays described by Vashistha and Naaman showed us what happens when polish is allowed to replace originality. Neither outcome is inevitable. Both can be resisted.

The task ahead is therefore not to insulate children from AI, nor to surrender to it, but to choreograph it so that the struggle remains theirs. Retrieval must still fire. Reflection must still sting. Feedback must still provoke. Practice must still be deliberate. Difficulty must still be desirable. These are the hard parts, and they cannot be outsourced without cost.

Education has always been about more than efficiency. It is about building minds capable of going beyond the information given, of speaking in voices not predicted, of caring enough to persist when the answer is not immediate. If AI is to have a place in that future, it must be bent towards those ends. Otherwise, we risk raising a generation fluent without memory, uniform without imagination, polished without pride.

Between silence and sameness lies the space we must defend: the messy, awkward, effortful work of thinking. That space is where learning happens. It is where originality is born. And it is what makes us, in the deepest sense, human.

Bibliography

Biesta, G. (2010) *Good Education in an Age of Measurement: Ethics, Politics, Democracy*. Boulder, CO: Paradigm Publishers.

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.

Bruner, J.S. (1973) *Beyond the Information Given: Studies in the Psychology of Knowing*. New York: Norton.

Brown, P.C., Roediger, H.L. and McDaniel, M.A. (2014) *Make It Stick: The Science of Successful Learning*. Cambridge, MA: Harvard University Press.

Carr, N. (2010) *The Shallows: How the Internet Is Changing the Way We Think, Read and Remember*. London: Atlantic Books.

Carless, D. (2015) *Excellence in University Assessment: Learning from Award-Winning Practice*. London: Routledge.

Cuban, L. (1986) *Teachers and Machines: The Classroom Use of Technology since 1920*. New York: Teachers College Press.

Dehaene, S. (2009) *Reading in the Brain: The New Science of How We Read*. New York: Penguin Viking.

Darling-Hammond, L. (2008) *Powerful Learning: What We Know About Teaching for Understanding*. San Francisco: Jossey-Bass.

Dweck, C. (2006) *Mindset: The New Psychology of Success*. New York: Random House.

Ericsson, K.A., Krampe, R.T. and Tesch-Römer, C. (1993) 'The role of deliberate practice in the acquisition of expert performance', *Psychological Review*, 100(3), pp. 363–406.

Flavell, J.H. (1979) 'Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry', *American Psychologist*, 34(10), pp. 906–911.

Gopnik, A. (2009) *The Philosophical Baby: What Children's Minds Tell Us About Truth, Love, and the Meaning of Life*. New York: Farrar, Straus and Giroux.

Haidt, J. (2023) *The Anxious Generation: How the Great Rewiring of Childhood is Causing an Epidemic of Mental Illness*. New York: Penguin Press.

Hattie, J. (2009) *Visible Learning: A Synthesis of Over 800 Meta-Analyses Relating to Achievement*. London: Routledge.

Immordino-Yang, M.H. (2016) *Emotions, Learning, and the Brain: Exploring the Educational Implications of Affective Neuroscience*. New York: W.W. Norton.

Kirschner, P.A., Sweller, J. and 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), pp. 75–86.

Laurillard, D. (2012) *Teaching as a Design Science: Building Pedagogical Patterns for Learning and Technology*. London: Routledge.

Livingstone, S. (2019) *Children: Risk and Safety on the Internet*. Cambridge: Polity Press.

Maes, P. (2023) *Commentary on MIT Media Lab research on AI and learning*. MIT Media Lab Reports. [Available at: <https://www.media.mit.edu>]

O'Neil, C. (2016) *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown Publishing.

Organisation for Economic Co-operation and Development (OECD) (2019) *OECD Learning Compass 2030*. Paris: OECD Publishing.

Posner, M.I. and Rothbart, M.K. (2007) *Educating the Human Brain*. Washington, DC: American Psychological Association.

Postman, N. (1985) *Amusing Ourselves to Death: Public Discourse in the Age of Show Business*. New York: Viking.

Roediger, H.L. and Karpicke, J.D. (2006) 'Test-enhanced learning: Taking memory tests improves long-term retention', *Psychological Science*, 17(3), pp. 249–255.

Sahlberg, P. (2011) *Finnish Lessons: What Can the World Learn from Educational Change in Finland?* New York: Teachers College Press.

Schleicher, A. (2018) *World Class: How to Build a 21st-Century School System*. Paris: OECD Publishing.

Schneider, S. (2024) *Artificial You: Child Development and Ethics in the Age of AI*. Princeton, NJ: Princeton University Press.

Sejnowski, T.J. (2023) *The Deep Learning Revolution*. Cambridge, MA: MIT Press.

Selwyn, N. (2016) *Education and Technology: Key Issues and Debates*. London: Bloomsbury Academic.

Sweller, J. (1988) 'Cognitive load during problem solving: Effects on learning', *Cognitive Science*, 12(2), pp. 257–285.

Turkle, S. (2011) *Alone Together: Why We Expect More from Technology and Less from Each Other*. New York: Basic Books.

Twenge, J. (2017) *iGen: Why Today's Super-Connected Kids Are Growing Up Less Rebellious, More Tolerant, Less Happy—and Completely Unprepared for Adulthood*. New York: Atria.

Twenge, J. (2024) *The Anxious Generation: How the Great Rewiring of Childhood is Causing an Epidemic of Mental Illness*. New York: Atria Books.

Vashistha, A. and Naaman, M. (2023) *Converging Voices: Homogenisation in AI-Supported Student Writing*. Cornell University arXiv preprint.

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*. San Francisco: Jossey-Bass.

Williamson, B. (2017) *Big Data in Education: The Digital Future of Learning, Policy and Practice*. London: Sage.

Wolf, M. (2018) *Reader, Come Home: The Reading Brain in a Digital World*. New York: HarperCollins.

World Economic Forum (2020) *Shaping the Future of Technology Governance: Artificial Intelligence and Machine Learning*. Geneva: World Economic Forum.

Zhao, Y. (2012) *World Class Learners: Educating Creative and Entrepreneurial Students*. Thousand Oaks, CA: Corwin Press.

Zuboff, S. (2019) The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. New York: PublicAffairs.

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 with 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.



For more information contact Dr Neil Hopkin at:

www.sunmarke.com

www.risdubai.com