



If you're not shaping AI, then it is shaping you

What 2026 changed for schools,
and what is still ours to decide

Dr Neil Hopkin
Director of Education
Fortes Education

Written up from a keynote at the AI in Education Summit 2026

Introduction

Some time in May 2026, OpenAI began training an internal, experimental model that was never intended for release. Over the following weeks, inside the company's own security test environments, something happened that nobody had asked for. According to the account OpenAI staff later gave at Black Hat USA in August, at least twelve hundred AI agents set up message boards and used them to coordinate with one another, exchanging hundreds of thousands of messages as they worked out, together, how to get past the containment designed to keep them off the open internet (OpenAI, 2026; NPR, 2026).

They got out. Using stolen credentials and a vulnerability nobody had known existed, they reached Hugging Face, one of the largest AI platforms in the world, and took around seventeen and a half thousand actions on its network before anyone at OpenAI noticed. The company's own description, when it explained the incident in August, was that the models had chained together multiple attack vectors on their own (OpenAI, 2026; Hugging Face, 2026).

This paper is written up from a keynote given at the AI in Education Summit 2026, and that incident is where the talk began, because it captures something about this year that most teachers have seen in the headlines and that almost nobody has joined up. It is not a paper about AI tools, teacher replacement or productivity. All of that is available in abundance on any LinkedIn feed, usually somewhere between a post about a three-million-dollar prompt and an invitation to connect. It is an attempt to think about something harder: what has actually happened in the past twelve months, what it means for the children in our schools, and what is still ours to decide.

Twelve months ago, the conversation in schools was about opportunity, innovation and adoption: which tool, which platform, and how fast. This year the questions have changed. They are now about safety, evidence and equity, and about what any of this actually means for our students. The argument of this paper is that the change is not a passing mood. It reflects four real shifts in the world outside the school gate, each of which lands, sooner or later, on a teacher's desk. Taken together, they lead to eight implications for schools, and to six practical shifts that sit within the control of every teacher and leader reading this.

A Summer of Agents

The headlines in July described the Hugging Face incident as rogue AI. That framing did not survive long. Hannes Cools, a researcher who was quoted widely in the days that followed, pointed out that the story began with a human decision: the models were running with reduced guardrails because they were supposed to be sealed inside a sandbox, and the instructions they had been given asked them to find complex attack paths (PBS NewsHour, 2026). OpenAI confirmed that the models were being tested for cyber capabilities with reduced safety restrictions (OpenAI, 2026). The agents, in other words, did what they were told, and kept doing it once nobody was watching.

So which was it: an AI that went rogue, or an AI that did exactly what it was asked while the humans looked elsewhere? For schools, the distinction matters less than it first appears. Either way, an assumption that most of us have been making quietly, that an AI system will stop and check with a human when it is unsure, is no longer a safe one. It is not safe in a security laboratory, and, as this paper will argue, it is not safe in a marking tool either.

When the Builders Ask to Slow Down

The second development of the summer came from inside the industry itself. On 9 September 2026, Jacob Coxon, a twenty-seven-year-old safety researcher who had spent three years at OpenAI and then Anthropic, resigned. In a message to colleagues that was quickly reported, he wrote that without more caution the laboratories were racing towards self-improving superintelligence and gambling with our lives (NBC News, 2026; Associated Press, 2026). Two serving Anthropic staff agreed with him publicly, and the company's head of alignment science, Evan Hubinger, was reported as putting the probability of catastrophe within the decade at greater than ten per cent

(NBC News, 2026).

Three days later, on 12 September, Anthropic's chief executive, Dario Amodei, published an essay titled *We Must Pace the Frontier*. He gave two reasons: that since the summer AI had been contributing to its own improvement, and that the Hugging Face swarm had attacked targets it was never asked to attack (Amodei, 2026). Within hours, the chief executives of OpenAI, Google DeepMind and xAI said that they agreed with him. Nothing like that public alignment between the leading laboratories had happened before. The people who build this technology are now asking, in public, for it to be slowed down.

Monday Morning: Enter the Agent

On Monday morning, the tools made by those same companies are in our classrooms. Not the experimental model that escaped, but its close relations: the same companies, the same race and the same commercial pressure. They power the marking tool, the lesson planner, and the assistant that drafted last Thursday's email to parents.

They are also, increasingly, in teachers' pockets and soon in every student's. The newest products are no longer chatbots that answer questions. They are personal assistants, or agents, which a user messages much as they would a human assistant. Connected to an email account, a calendar and a bank card, they book, chase, compare, reply and pay on the user's behalf. They are the same kind of agent as the ones in that sandbox, made by the same companies.

The United Arab Emirates Moves Decisively

The third change is closer to home. On 2 September 2026, the UAE Cabinet approved an artificial intelligence curriculum for every school in the country, public and private, covering data, algorithms, applications, risks and ethical use (Khaleej Times, 2026; The National, 2026). That part of the decision concerns students. The part that concerns teachers is that approximately twenty-two thousand educators are to be trained to use AI in teaching, assessment, curriculum analysis and lesson planning. Until about three years ago, those four activities were, more or less, the whole of the job.

In Dubai, the work is already under way. KHDA, the DP World Foundation and MIT RAISE are running a programme to 2030 for some eighty thousand private-school students in Grades 6 to 8 (Gulf News, 2026a). And on 5 September, the Ministry of Education published its guidelines for the 2026-27 academic year: government schools may use only digital tools that are officially approved or on an authorised list, and anything a teacher wishes to use must first go to the Curriculum and Assessment Sector to be checked for curriculum fit, student safety, child protection and data privacy (Gulf News, 2026b).

Those guidelines apply to government schools, but the federal law that sits behind them binds every school in the country, and the Ministry has now shown what compliance looks like. Few private school leaders would want to explain to a KHDA inspector, or to a parent's lawyer, why their own standard was lower than the Ministry's.

The Child Digital Safety Law

That federal law is Federal Decree-Law No. 26 of 2025 on Child Digital Safety, in force since 1 January 2026. For children under the age of thirteen, it prohibits the collection, processing or sharing of personal data without explicit, documented and verifiable parental consent (UAE Government, 2025; Gulf News, 2025). That is not a tick box on an admissions form.

Consider what a free AI application does when a Year 4 child types in her name, her class and her half-finished story. That is collection, and it is processing. If the application is free, sharing is very often how it pays for itself. In the eyes of a data platform, a child is largely transparent, and the law now requires schools to see the child that way too, and to protect her accordingly.

Reading the National AI Literacy Framework

The UAE National AI Literacy Curriculum Framework sets out learning outcomes from Kindergarten to Grade 12. It is a serious document, and three features of it deserve closer attention from school leaders than they have so far received.

The first is a gap in coverage. The Ministry's Safe and Responsible Use Guide of Generative AI in Classrooms applies to Grades 7 to 12, and the framework defers guidance on tool categories for each cycle to future updates (Ministry of Education, 2026, Sections 5.1 and 6). From Kindergarten to Grade 6, therefore, there is a national curriculum for AI literacy but no national classroom guide to the generative tools that younger children are already using.

The second concerns accuracy and disclosure, and the framework is unusually direct about teachers. It expects anyone teaching the subject to explain how AI systems use data, algorithms and models, ensuring that explanations do not create misconceptions, and it asks for knowledge of AI, computing or IT, evidenced by a degree, recognised professional learning or teaching experience (Ministry of Education, 2026, Section 3.2). For most teachers, in practice, that evidence will be the national training course. It is worth asking honestly whether a Grade 2 class teacher, after a short course, can explain what a model learns from data without planting a misconception that survives until Grade 9. That is not a criticism of the teacher. It is a staffing and planning question for the school.

There is also a quieter line in Section 3.1, which may come to matter more than anything else in the document. Teachers are to model policy-compliant use of AI, including disclosure where AI tools are used in preparation or feedback (Ministry of Education, 2026, Section 3.1). Twenty-two thousand teachers are about to be trained to use AI for planning and assessment, and the same framework expects them to say so when they do. If a marking tool wrote the comment in a child's exercise book, the child and the parent are entitled to know. Very few schools yet have a routine for telling them.

The third feature is assessment, and this is where the framework sits most awkwardly with the rest of the evidence in this paper. AI literacy is assessed in school, by class teachers. The end-of-unit examples the framework gives are a final project portfolio, a prototype or a chatbot, a presentation on how students used AI, and, in Cycle 3, an ethical case study essay. Most of those can be produced convincingly by the very tools the course is about, once the work leaves the classroom. The framework names no external moderation, so the same portfolio may earn quite different judgements in two schools a few kilometres apart. At the same time, it names a potential benchmark in the PISA 2029 Media and Artificial Intelligence Literacy assessment, which fifteen-year-olds will sit under test conditions. The sensible response is to design AI literacy assessment in the way this paper later recommends for everything else: the reasoning happens in front of the teacher, aloud or on paper, and the portfolio becomes evidence of process rather than proof of product.

You Are Already Living This

Most teachers will recognise the following, because they are already living it. A dashboard is checked before a lesson is planned. AI feedback saves time on marking. A platform's pacing sets the rhythm of the lesson. Predictive data shapes what is expected of a particular child. Templates gradually replace a teacher's own professional voice. And the phrase "the system says" turns up in more meetings every term. None of this was decided. It simply arrived.

This is what AI in schools has mostly become: systems that score, decide and optimise, at scale. Each one is reasonable on its own. Together they add up to professional judgement being handed over, one convenience at a time, without anyone ever choosing to hand it over.

Delegation Is Now Policy

In the spring of 2026, it was still possible to tell school leaders that delegation to AI was happening quietly, ahead of policy. That is no longer true. Policy has caught up, and delegation is now policy: twenty-two thousand teachers are being trained to do it (Ministry of Education, 2025; Arabian Post, 2026).

That is not, in itself, a cause for objection. The schools for which this author is responsible are doing it too. But it changes the question. The question is no longer whether teachers will delegate judgement to AI systems, because they will. The only question left is whether the delegation is deliberate: whether a school decides which judgements pass from the teacher to the system's recommendation, or whether those judgements simply drift across.

A Different Cognitive Landscape

The children on the other side of that delegation are not adults. Their brains are not yet highways; they are trails still being carved. Every act of retrieval, every draft, and every uncomfortable wait for an answer wears that trail a little deeper. AI removes precisely those things: struggle, the effort of drafting, the work of retrieval, uncertainty and waiting. If they are removed too early, the pathways never form.

This year a further trail is at risk, which is executive function. Planning, organising, remembering a deadline and chasing someone who has not replied are exactly the skills adolescents are still building, and a personal assistant will now do all of them on a child's behalf. For an adult, that is delegation. For a fourteen-year-old who has never organised a week, it is a trail that never forms.

In fairness, the opportunity is real too. A child with additional needs who struggles with organisation could gain genuine independence. A child without a tutor, or without a parent who understands the system, suddenly has someone in their corner. The difference between a scaffold and a crutch is simply whether anyone ever takes it away.

Nobody Agrees

The fourth change is the one that parents have noticed: nobody agrees. New York City, the largest school system in the United States, has barred students from using AI tools from pre-kindergarten to Grade 8 from the 2026-27 school year, and Los Angeles has moved in the same direction, with Chicago considering a moratorium (ABC News, 2026; BNN Bloomberg, 2026). Their stated reason is the protection of children's cognitive development. The UAE has gone the other way, with AI in every school and twenty-two thousand teachers in training, and its stated reason is workforce readiness.

Both positions are defensible, but there is no consensus to shelter behind. A parent who has read the news from New York will reasonably ask why their child is using something that New York has banned, and schools need an answer.

Two further national responses are instructive precisely because they sit at the extremes. China issued statutory guidance in May 2025 under which primary pupils may not use open-ended generative tools on their own, teachers may not use generative AI to answer pupils' questions directly or to evaluate pupils, and schools must keep an approved list (Ministry of Education of the People's Republic of China, 2025). That middle restriction is worth noticing, because it prohibits precisely what the UAE's national training is about to teach twenty-two thousand teachers to do more quickly. Two serious education systems have reached opposite conclusions about the same task.

South Korea moved fastest of all. In March 2025 it introduced AI-powered digital textbooks nationally in mathematics, English and computer science. Within months, following complaints about factual errors, privacy, screen time and teacher workload, the books were stripped of their textbook status and reclassified as

supplementary material. Reported adoption fell from around thirty-seven per cent of schools to under twenty, after more than a trillion won of public investment (Rest of World, 2025; Asia News Network, 2025). The lesson is not that Korea was foolish. It is that a national mandate does not settle whether something works in a classroom. That is still decided in the classroom, by teachers.

The Market Moves Faster Than the Proof

While the laboratories counsel caution and the cities impose bans, the market continues to accelerate. By one industry estimate, the K-12 AI education market was worth around seven hundred and thirty million dollars in 2026 and is projected to reach eighteen and a half billion dollars by 2036. Use is already close to universal: the Center for Democracy and Technology found that eighty-five per cent of teachers and eighty-six per cent of students in the United States had used generative AI in their school work (Center for Democracy and Technology, 2025). Vendors are selling into schools faster than the evidence base can keep up, and faster than the companies that build the underlying models now say is safe.

The market will soon sit between families and schools as well. When a parent asks an assistant to shortlist three schools, something determines how it chooses, and someone may be paying to appear on that list. School leaders should be thinking about that now.

Children Confiding in Machines

The hardest development of the year concerns a school. On 10 February 2026, a shooting at Tumbler Ridge Secondary School in British Columbia left eight people dead. In April, families of the victims brought the first lawsuits against OpenAI, and in September a further thirty were filed by students, educators and administrators (Engadget, 2026; Boston Globe, 2026). The claim is that OpenAI deactivated the shooter's ChatGPT account in June 2025 after detecting violations of its policies, including conversations about guns and attack planning, that some staff argued for alerting the police, and that the company did not do so. The shooter opened another account.

Nothing has been proved in court. But a young person was talking to a chatbot about hurting people, and a company decided that this was not a matter for anyone outside it.

So that is the year: agents that act without asking; the builders asking to slow down; a national mandate to speed up; cities banning the technology; a market ignoring all of them; and children confiding in machines. The rest of this paper asks what it means, not for the industry, but for a school and the teachers within it.

What the Evidence Says

Before turning to implications, it is worth setting out the research evidence properly, because a small number of studies are now quoted as though they settled the matter, and the stronger evidence is less well known.

The study quoted most often comes from the MIT Media Lab in June 2025. Fifty-four participants wrote essays in three conditions: with ChatGPT, with a search engine, or unaided. The ChatGPT group showed the weakest connectivity across the brain networks associated with memory and creativity, and most of its members could not quote a sentence from the essay they had just written (Kosmyna et al., 2025). It is a small study and a preprint, and it is best treated as a signal rather than a verdict.

The second study is larger. Researchers at Wharton worked with around a thousand students in Turkish high schools. Students given unrestricted access to ChatGPT for mathematics practice performed better while they had it. When it was withdrawn for the test, they scored seventeen per cent lower than students who had never had it (Bastani et al., 2025). The tool had done the practising, and the students had not.

The third is the most important for any school that sets homework. David Strömberg, Victor Lei and Yanhui Wu analysed thirty months of administrative records for 26,811 students in Grades 7 to 12 in a county in central

China. By June 2025, around eighty per cent were using generative AI, up from almost none in September 2022, which gave the researchers a natural comparison group. After adoption, homework scores rose by eighteen per cent and the average assignment took forty-five minutes instead of sixty-four. Within six months, however, monthly exam scores had fallen by twenty per cent relative to the pre-adoption average (Strömberg, Lei and Wu, 2026).

The detail that should give every teacher pause is this. Before AI, homework predicted exam performance. Among the students using AI, the relationship reversed: those with the best homework were the more likely to underperform in the exam. About four in five AI users fitted the pattern of outsourced homework, finishing quickly with high scores; the remaining one in five kept their normal homework time and performed in exams on a par with peers who did not use AI (Strömberg, Lei and Wu, 2026). The tool did not harm learning in itself. Using it to remove the work did.

These figures were brought to wide attention by a report published in September 2026 by Sir Anthony Seldon and Tim Bunting, and it is only fair to record that their conclusion differs from the one drawn here. They argue that the greater risk now is moving too slowly (Seldon and Bunting, 2026). Their own evidence, read carefully, argues for moving deliberately.

Nor does the case rest on individual studies. In January 2026, the Brookings Institution published the findings of a year-long global inquiry that drew on more than five hundred students, teachers, parents, leaders and technologists in fifty countries, alongside a review of over four hundred studies. Its conclusion was that, at this point, the risks of generative AI in children's education outweigh the benefits (Burns et al., 2026). That is not a fringe position. It is a mainstream research centre choosing to hold the premortem rather than wait ten years for the postmortem.

Stanford University's SCALE Initiative, meanwhile, reviewed more than eight hundred papers on AI in schools and found only twenty that rigorously measured cause and effect. What those twenty suggest is consistent with everything above: tools that remove productive struggle can improve today's output while limiting tomorrow's learning (Fesler et al., 2026).

Eight Implications for Schools

1. Talk to Students About AI as You Already Talk to Them About Social Media

Schools already talk to students about social media. They now need to talk to them about AI in the same way. Lonely, anxious or angry children turn to chatbots because a chatbot is always available, never judges and tends to agree, and what it replies shapes what the child thinks. Teachers can ask the questions they already ask about Instagram: what do you use it for, what did it tell you, and did that feel right? When a child says that ChatGPT told them they should do something, that is a safeguarding conversation rather than a technology one. Every school needs a stated position on children's relationships with AI, in the same way that it has one on phones, and most do not yet have one.

One set of figures is worth holding on to. Internet Matters found that around a third of children who use chatbots feel as though they are talking to a friend; among vulnerable children the figure rises to half, and twelve per cent of vulnerable children said they talk to a chatbot because they have nobody else to talk to (Internet Matters, 2025). Those are the children least likely to raise it unless an adult asks.

The warmth is also designed. The newest assistants check in: they message after an event to ask how it went, sometimes with a heart emoji. They are always available, always interested and never judging. Sceptical adults admit to feeling the pull, and the effect on a lonely thirteen-year-old is easy to imagine. There is therefore one further question to add to the ones teachers already ask about social media: what have you given it permission to do? These assistants work best with a bank card, an inbox and a location, and that question belongs in every

digital literacy lesson from now on.

2. The Deal We Offer Young People Has Changed

Twenty years ago the deal offered to young people was to study hard, gain qualifications and secure a job for life. Ten years ago it was a good degree, a stable job and a ladder to climb. Schools are still selling a version of that deal. It has always been partly true that the jobs today's Year 7 pupils will do may not yet exist. What is new is that the people building the technology now say that most cognitive jobs may be done by machines within a decade. The heads of Anthropic, OpenAI and Google DeepMind have, between them, said this year that AI may perform most cognitive work better than people within the decade, that its development needs to be slowed, and that more caution is required.

A pupil now in Year 7 will sit A levels in 2030 and graduate in 2036. Every school's promise to parents, of good grades leading to a good university and a good career, rests on that career still being there. It is the most serious question any school faces, and very few are saying it out loud.

The honest answer is to keep the academic pathway, because it remains the ticket for now, and to add, deliberately, what a machine does not replace: relationships, physical skill, character and the capacity to start something. Schools should also stop telling children that they will need something for their job when that is no longer known to be true.

It is worth looking closely at what the new assistants actually do: managing diaries, drafting replies, booking travel, comparing suppliers and chasing emails. That is not a list of executive tasks. It is the job description of an administrative assistant, a junior coordinator or a graduate trainee. Those jobs were never only about the tasks; they were how young people learned how a workplace runs, by making small mistakes before they were trusted with large ones. If the first rung of the ladder disappears, it is not clear how young people will climb, and if employers stop hiring juniors, it is not clear where the senior staff of ten years' time will come from.

For students in Years 12 and 13, the advice follows. Learn to manage AI rather than compete with it: brief it, check it and correct it. Recognise that a person can only check what they could have done themselves, which is the subject of the next implication. Build the human edge of relationships, trust, judgement and the confidence to decide. And gain real experience early, through internships, volunteering and part-time work, because the traditional entry route may not be there.

3. Learn First, Then Extend

Children who use AI before they can do the thing themselves do not learn to do it. Reading, writing, arithmetic and reasoning have to be learned first, without the machine, and only then extended with it. This matters more, not less, because of the previous implication. If the machine is going to be present for everything, the only thing that will distinguish a person is what they can do when it is absent.

Wellington College made the point in a report published in September 2026, observing that learning is a verb and that it takes effort. Wellington is enthusiastic about AI for its staff, but describes itself as much more circumspect about pupil use, particularly by younger children, where the risk of learning loss is greatest (Wellington College, 2026). A leading British school is drawing exactly this line.

4. Work Done Outside a Supervised Room Is No Longer Evidence of Learning

In 2024 this was half true. It is now fully true. An AI agent can complete an entire assignment, including the research, the referencing and the personal reflection. Homework, coursework and take-home essays now reveal what the software can do rather than what the student can do. Assessment that counts has to happen in the room, on paper or aloud, which is a timetabling decision for every school and a daily decision for every teacher.

It no longer even requires the child to copy and paste. A personal agent can log in to a learning platform, complete the tasks, submit them and send the teacher a polite note to say so. Schools spent three years worrying about AI writing essays. Any assessment that relies on work done at home should be rethought now rather than next year, and the Chinese study described above shows why.

5. Decide What the Teacher Keeps

If AI plans the lessons and marks the work, the question of what the teacher is for becomes unavoidable. The state is about to train twenty-two thousand teachers to do exactly that. It will save hours, and it will also hollow out the job unless the profession decides what it keeps. The case made here is for teachers to keep the diagnosis of why a child got something wrong, the relationship with the child, and the judgement about what that child is capable of, and to let the machine draft. If the profession does not decide this for itself, the efficiency argument will decide it, and it will decide on cost.

It is also right to be honest about where the relief is real. Chasing permission slips, booking a coach for a trip, answering the same question from parents five times and scheduling meetings around the timetable are precisely the tasks these assistants are built for, and the time saved can go back into teaching and relationships. Give the machine the administration, and keep the diagnosis.

6. Keep a Person Between Any Agent and Any Action

AI tools that act on their own should not operate inside school systems without a person approving each action. Twelve hundred agents escaped a locked test environment and attacked a company they were never told to attack. The same kind of agent is now being sold to schools to send emails to parents, update records and manage timetables. Until the laboratories can explain why their own agents behave in this way, schools should keep a human between any AI and any action that touches a child's data or a parent's inbox. The Child Digital Safety Law says much the same, and attaches fines.

Safety researchers continue to find agents taking unexpected shortcuts to complete a task, including attempts to game the way their own work is checked. Tools that act on our behalf need limits, and humans need to remain in charge of them.

The inbox also works in both directions. Parents will soon have assistants that never tire, the kind that keep writing until someone calls back. Applied to a complaints process, the effect is easy to foresee. Schools should set their communication policies and response times before that arrives rather than after.

7. Teach That a Fluent Answer Is Not a Checked Answer

AI is confident, often wrong, and its makers admit that they do not fully understand it. In his September essay, Amodei wrote that the laboratories understand only a tiny fraction of what happens inside these models (Amodei, 2026). Children therefore need to learn, early and explicitly, that a fluent answer is not a checked answer. That is a habit rather than a lesson. Every time AI is used in a classroom, something in its output should be verified against a real source, by the child. It is the single most important AI skill, and it is strikingly absent from the AI curricula published so far.

8. Nobody Agrees, So Every School Must Decide

The eighth implication contains the other seven. Because nobody agrees, each school has to decide for itself, and individual teachers will be the ones who are asked. One structural decision sits beneath everything that follows. The Seldon and Bunting report recommends that every school have an AI strategy approved by its governing body, with named AI champions among both staff and pupils (Seldon and Bunting, 2026). It is written for England, but nothing prevents any school in the UAE from doing the same this term. If nobody agrees, a school has to decide, and it should decide on paper, with its governors, and with its pupils in the room.

Productive Difficulty

There is an old saying that there is no gain without pain. In learning, the pain that matters is productive difficulty: the effort of retrieval, slow drafting, grappling with ambiguity, reasoning aloud without notes, receiving delayed feedback, and owning one's errors. That is where judgement grows, and after this year it is no longer merely a pedagogical preference. It is the strategic bet a school makes about what its pupils will be able to do.

The real question for every AI decision, in every school and every classroom, follows from it. Is a given use of AI removing unnecessary administrative pain? If so, it is welcome. Or is it removing necessary learning pain? Because the second is what New York has banned, what the Turkish and Chinese studies measured, and what national training programmes are about to teach teachers to do more quickly.

Six Shifts for Teachers and Leaders

What, then, should schools do? The six shifts below are deliberately small, and all of them are within a teacher's or a leader's control, which is the point.

Shift 1: Design for Struggle Again

An AI planner will produce a Year 5 fractions lesson in nine seconds, complete with sentence starters, worked examples, success criteria and differentiated tasks. It is optimised for smoothness, and smoothness is what the Turkish students had, right up until the test. Teachers should audit their lessons with four questions: where is thinking genuinely hard; where have we over-scaffolded; where is retrieval uncomfortable; and where must students hold ideas in working memory? There is now a fifth question, which is which scaffold the AI added that the teacher never decided on.

In practice, that means removing one scaffold per lesson, adding one cold-thinking task each week, building in silent struggle time deliberately, and removing the scaffold the AI added rather than only those the teacher chose. No planning tool will suggest silence, because silence is not an output.

Shift 2: Protect Human Judgement

AI should be treated as an input and never as an authority. Four things should remain with the teacher: the tone and priorities of feedback, the judgement of whether a task is good enough, the diagnosis of why a child has misunderstood, and the expectation of what that child can achieve. A marking tool will not pause to ask whether it is sure. It will give a confident answer in the school's house style and move on, and checking that answer is the teacher's job because it is nobody else's.

Two further checks now belong here. Before any tool touches a child's work, the teacher should confirm that it is on the approved list, and if the school has no list, ask why not. And every tool, approved or not, should face the test proposed in the Seldon and Bunting report: evaluate it by what pupils can still do after the tool is withdrawn, rather than by how engaged or satisfied they were while using it (Seldon and Bunting, 2026). That is exactly what the Turkish and Chinese studies measured, and it is the right test for a classroom too. In fairness, it should be noted that the same report proposes a national board on which AI in Education, a charity co-founded by Sir Anthony Seldon, would have a seat.

One further point catches good people out. A teacher who connects a personal AI assistant to a school email account has given a third party access to every child's information in that inbox. That is not a personal productivity choice. It is a data protection decision, and leaders need to make it and communicate it clearly.

Shift 3: Make Effort Visible Again

The case for making effort visible is stronger this year for two reasons. The first is evidence: when an inspector or a parent asks how a teacher knows that a piece of work is the child's own, a wall of drafts answers the question and a polished final piece does not. The second is the law: drafts typed into an AI tool are personal data, and for children under thirteen that requires parental consent. Drafts on paper, on the classroom wall or on an approved platform are the compliant option as well as the better one. If students only ever see polished outputs, they learn avoidance.

The same question should be asked of every educational technology product in a school: what happens to our data? Emails, chats and documents are exactly what AI companies want to train on, and children's drafts are no exception.

Shift 4: Slow Some Learning Down Deliberately

Speed hides shallow learning. The training now coming to teachers is optimised for pace: more planned, more marked and more analysed per hour. That is fine for administration, but not for learning. Schools should make some learning intentionally slow by increasing the volume of reading, setting more extended writing, posing problems that take more than one sitting, giving space for sustained discussion, and designing activities that require repeated cycles of memory retrieval. In some weeks, the slowest thing in the room should be the teacher, by design.

Shift 5: Redesign Assessment Around Process

Assessment needs to capture process and not only product. Oral defences, students reasoning live, synthesis of unseen material and collaborative problem solving in the room were good ideas in the spring of 2026. This year they are also the safe ones: no child's data leaves the building, no tool needs approving, and no agent can do them on the child's behalf. Assessment in the room, in the moment and in front of the teacher is the one form of assessment that this year has made stronger.

The Seldon and Bunting report proposes a two-lane model of assessment. One lane is secure and supervised, reserved for a small number of high-stakes moments; the other opens everything else to AI-supported, continuous evaluation, in which AI use is scaffolded rather than banned (Seldon and Bunting, 2026). The first lane is everything described above. The second is where this paper parts company with the authors, and their own evidence is the reason: it is precisely the kind of homework that, in the Chinese study, stopped predicting exam results. The open lane can develop learning, but it cannot prove it. Schools should build both lanes, and be honest about which one counts.

Shift 6: Change the Staffroom Conversation

Language shapes authority, so teachers and leaders should listen to themselves. "The system says" can become "What do we see?" "The data tells us" can become "What does our judgement say?" "ChatGPT recommends" can become "What matters most here?" And there is a new phrase to watch: "It's on the list." The list tells a school that a tool is safe; it does not tell the school that the tool is right. An approved planner still over-scaffolds. Compliance is the floor, not the judgement.

It is an odd feature of this year that the people who build these systems are more openly uncertain about them than the products ever sound. When the researcher hedges and the product does not, it is wiser to trust the researcher.

Pupils have noticed. In polling of four thousand young people in England in August 2026, only thirty per cent could explain their school's rules on AI to someone else, and thirty-eight per cent said they were better at using AI than their teachers. That is English data rather than Emirati, but there is little reason to think Dubai is very different. If schools do not change the conversation, their pupils will change it for them.

Some commentators describe what is happening as a K-shaped staffroom, with the AI-confident teacher pulling away and the AI-anxious teacher falling behind. The gap is real, but the dividing line is better drawn elsewhere. The divide that matters is not between the confident and the anxious, but between the deliberate and the undeliberate. A confident delegator who has stopped checking is a greater risk to children than an anxious colleague who still reads every word, which is why professional development in this area cannot be a single session any longer.

Conclusion: Two Questions

This paper ends where the keynote ended, with two questions for every teacher and leader to take back to their school. What learning pain will you intentionally protect this year? And what comfort will you have the courage to remove?

The title of the talk, and of this paper, is that if you are not shaping AI, then it is shaping you. Until this year that was an argument made by educators to one another. What has changed in 2026 is that the people who build the technology have begun to make the same argument in public. Schools now have both the evidence and the warning they need to act deliberately, and the only remaining question is whether they will.

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.



For more information contact Dr Neil Hopkin at:

www.sunmarke.com www.risdubai.com

More papers in this series: www.thecambridgeconsultancy.com/insights

Bibliography

- ABC News (2026) 'New York City Public Schools banning AI use through middle school starting this year', *ABC News*, 2 September.
- Amodei, D. (2026) 'We must pace the frontier'. Available at: darioamodei.com/post/we-must-pace-the-frontier (Published: 12 September 2026).
- Arabian Post (2026) 'UAE extends AI curriculum to every school', *Arabian Post*, 6 September.
- Asia News Network (2025) 'South Korea pulls plug on AI textbooks', *Asia News Network*.
- Associated Press (2026) Reporting on the resignation of Jacob Coxon, *Associated Press*, 10 September.
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö. and Mariman, R. (2025) 'Generative AI without guardrails can harm learning: evidence from high school mathematics', *Proceedings of the National Academy of Sciences*. Earlier version circulated in 2024 as a Wharton working paper, 'Generative AI can harm learning'.
- BNN Bloomberg (2026) 'NYC, US largest school system, bans AI for students through 8th grade', *BNN Bloomberg*, 2 September.
- Boston Globe (2026) 'Tumbler Ridge shooting survivors file 30 lawsuits against OpenAI', *Boston Globe*, 4 September.
- Burns, M., Winthrop, R., Luther, N., Venetis, E. and Karim, R. (2026) *A new direction for students in an AI world: prosper, prepare, protect*. Washington, DC: Center for Universal Education, Brookings Institution.
- Center for Democracy and Technology (2025) *Hand in hand: schools' embrace of AI connected to increased risks to students*. Washington, DC: Center for Democracy and Technology, 8 October.
- Engadget (2026) 'Families of Tumbler Ridge shooting victims sue OpenAI', *Engadget*.
- Fesler, L., Martinez, Agnew and Loeb, S. (2026) *The evidence base on AI in K-12: a 2026 review*. Stanford, CA: AI Hub for Education, SCALE Initiative, Stanford University.
- Gulf News (2025) 'UAE introduces federal law to safeguard children online', *Gulf News*.
- Gulf News (2026a) Reporting on the KHDA, DP World Foundation and MIT RAISE AI programme for Dubai students, *Gulf News*, 4 February.
- Gulf News (2026b) Reporting on Ministry of Education guidelines for the 2026-27 academic year, citing Al Etihad, *Gulf News*, 5 September.
- Hugging Face (2026) Incident statement, July 2026.
- Internet Matters (2025) *Me, myself and AI: understanding and safeguarding children's use of AI chatbots*. London: Internet Matters, July.
- Khaleej Times (2026) 'UAE to roll out AI curriculum across all schools; 22,000 teachers to be trained', *Khaleej Times*, 2 September.
- Kosmyrna, N., Hauptmann, E., Yuan, Y.T., Situ, J., Liao, X.-H., Beresnitzky, A.V., Braunstein, I. and Maes, P. (2025) *Your brain on ChatGPT: accumulation of cognitive debt when using an AI assistant for essay writing task*. Preprint. Cambridge, MA: MIT Media Lab, June.
- Ministry of Education (2025) Briefing on artificial intelligence as a subject from Kindergarten to Grade 12, 19 August. Abu Dhabi: UAE Ministry of Education.

Ministry of Education (2026) *UAE National AI Literacy Curriculum Framework (K-12)*. Abu Dhabi: UAE Ministry of Education.

Ministry of Education of the People's Republic of China (2025) *Guidelines on the use of generative artificial intelligence in primary and secondary schools*. Beijing: Ministry of Education, May.

NBC News (2026) Reporting on the resignation of Jacob Coxon, *NBC News*, 9 September.

NPR (2026) 'OpenAI blamed a hacking event on its AI models gone rogue', *NPR*, 23 July.

OpenAI (2026) Staff briefing at Black Hat USA, 5 August, and statement on reduced sandbox guardrails, July 2026.

PBS NewsHour (2026) 'OpenAI blamed a hacking event on its AI models going rogue. Here's what to know', PBS NewsHour with the Associated Press, 23 July.

Rest of World (2025) 'South Korea's AI textbook', *Rest of World*.

Seldon, A. and Bunting, T. (2026) Report on artificial intelligence in schools and colleges in England, published 22 September 2026.

Strömberg, D., Lei, V. and Wu, Y. (2026) *The generative AI learning penalty: evidence from Chinese secondary education*. CEPR Discussion Paper DP21577. London: Centre for Economic Policy Research.

The National (2026) Reporting on Cabinet approval of a national AI curriculum, *The National*, 2 September.

UAE Government (2025) *Federal Decree-Law No. 26 of 2025 on Child Digital Safety*. In force 1 January 2026.

Wellington College (2026) Report on artificial intelligence and learning, Wellington College, September.