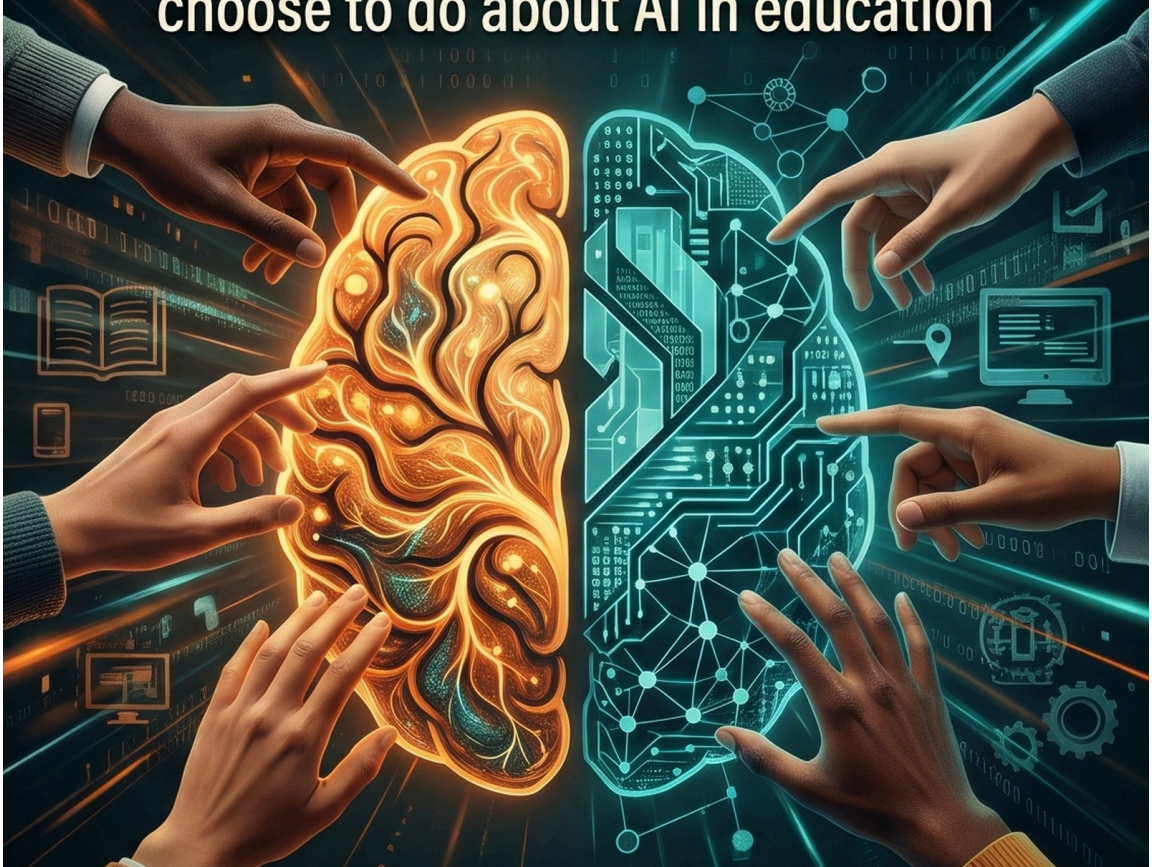
DELIBERATELY HUMAN

What schools, families and governments choose to do about AI in education



Paul Penfold

A groundbreaking look at choices, governance, and the future of learning

Deliberately Human

Deliberately Human

What schools, families and governments choose to do about AI in education

Paul Penfold

Deliberately Human

What schools, families and governments choose to do about AI in education

© 2026 Paul Penfold. Some rights reserved.

This book is licensed under the Creative Commons Attribution–NonCommercial–NoDerivatives 4.0 International (CC BY-NC-ND 4.0). You are free to copy and share it in any medium or format, as long as you credit the author, do not use it for commercial purposes, and do not remix, adapt or build on it. The full licence is at https://creativecommons.org/licenses/by-nc-nd/4.0/.

In plain words: free to use and share, but not to rework, and not to sell or use commercially. Short quotations with credit are always welcome. For anything beyond this licence, please ask the author.

The licence covers the text and the original figures. Material quoted from other sources stays under its owners' terms.

ISBN (print): 978-1-0655518-0-5 ISBN (EPUB): 978-1-0655518-1-2

How to cite: Penfold, P. (2026). Deliberately Human: What schools, families and governments choose to do about AI in education (First edition, October 2026). https://deliberately-human.netlify.app/. To cite one of the underlying papers, use its DOI from the list of original papers. Please give the edition and the date you read it, because this book is updated.

The Introduction and Part Eleven explain how this book was made, who did what, and what checking found. The papers behind it appear in Proceedings of Future Science Archives, 2026.

First edition, October 2026.

For my mother, a free spirit, who inspired me. I have followed in her footsteps.

Acknowledgements

A few people helped me change direction and find the work I love most.

At Hong Kong Polytechnic University, Professor Kaye Chon and Professor Haiyan Song believed in me, took me into their faculty and gave me the freedom to develop resources, tools and ideas. Those included the Second Life virtual campus and virtual hotel. Josie Szeto and Peter Duffy, of the university's Educational Development Centre, backed my ideas and helped me find funding to build tools and apps for the university.

After I left the university in 2010, I moved into education consultancy. Eva Raute of GOPA Consulting and Véronique Meyers-Sotil of Luxembourg Development inspired and supported me there.

Here in Chiang Mai, I am grateful for the encouragement and example of John Ho (Alt Chiang Mai), Peter Richards (Harmony Hualin) and Martin Venzky-Stalling (Chiang Mai University and TEDEd).

I am also grateful to the many expert researchers and writers on LinkedIn whose work I value and have learned from.

My thanks also to Richard L. Wilson, who read the research early and wrote the review that opens this book.

The eleven papers behind this book are published openly in Future Science Archives and archived on Zenodo. I am grateful to both for making open, citable publication possible for applied education research.

This book was written with the help of AI tools, which drafted, checked and built much of it under my direction. Part Eleven explains who did

what and what the checking found. Any errors that remain are mine.

“Should we allow AI?” is the wrong question.

Schools are being asked to decide about AI faster than the evidence arrives. This book argues that the choice is not about what the technology can do. It is about what schools, families and governments decide to do with it.

Drawing on a series of linked research papers, it looks at learning, children’s wellbeing, AI literacy, assessment, entry-level work, teachers, leadership, governance and equity. Every claim is labelled by strength, from established trend to expert projection to speculation, so you can see which parts deserve trust and which deserve scepticism.

It is written for school leaders, teachers, parents and policymakers. The online edition adds interactive tools, a coach for each part, and a search that answers from the book’s own text.

Ask instead: what kind of AI, used how, and checked by whom?

Reviews

"Paul, this is excellent; really!

I was drawn in by this research. What struck me first is how honest it is. It is honest, it feels honest, and it reads honest. Labelling every claim, from established trend to speculation, is something I wish more writing about AI in education did. It lets the reader trust the parts that deserve trust and think critically about the rest.

For me, the most immediately useful message for school leaders is that how AI is designed matters more than whether it is used. The study you cite says it all: an unrestricted chatbot raised practice scores but lowered exam results, while a tutor-style version that made students think did not. Many schools are still asking "Should we allow AI?" Your article moves the question to "What kind of AI, used how, and checked by whom?" That is a question a principal can act on tomorrow, in the next procurement decision or the next assessment policy.

I also found the "moves that hold in every future" very practical, especially teaching verification and assessing what learners can do without the tool. Those connect closely with my own work on teaching students how to learn, not just what to learn."

Richard L. Wilson, M.Ed.
Educator | Writer | Editor | Academic Director

Contents

Foreword 12
Introduction 14

Part One: Choices, Not Capability **18**
1. The class of 2034 20
2. Seven pressures already in the room 22
3. What other scenario sets say, and where they disagree 25
4. Four futures for 2035 27
5. Three futures for 2045 29
6. Six switches and what to do now 31

Part Two: Learn First, Then Use AI **37**
7. The wrong question: does AI help or harm learning? 39
8. What thirty experts found 41
9. What learning science would predict 44
10. Four studies tested against the consensus 46
11. Where the evidence and the consensus part 49
12. Learn first, then use AI: the rule and what it means for schools 51

Part Three: Growing Up with Machines That Talk Back **56**
13. The student at lunch: a question of care 58
14. Relationships are how learning works 59
15. Effort, pace and the habit of checking 62
16. What children are owed: the truth about what they talk to 64
17. Responses in 2026 and five developmental safeguards 66

Part Four: What Counts as AI Literacy **71**
18. Everyone agrees it matters, nobody agrees what it is 73
19. Five frameworks side by side 76
20. Four points of agreement, five real differences 80
21. China, Boston and New York City 84
22. What the definitions leave out, and a model for schools 87

Part Five: When the Artefact Is No Longer Evidence 93
23. The quiet assumption behind most assessment 95
24. How big is the problem? 97
25. Why detection is the wrong answer 100
26. From integrity to validity 102
27. What schools can build, and the risks of building it 105

Part Six: The Broken Bottom Rung 113
28. The place where people learn to work 115
29. What the labour market is signalling 117
30. The expertise paradox 122
31. Credentials, caring work and a three-layer model 124

Part Seven: Teachers at the Centre 129
32. A mismatch at the heart of school AI 131
33. How big is the gap? 133
34. What AI can offer teachers, and what the time is for 137
35. Professional learning that works 140
36. Agency, the right to say no, and a capability pathway 143

Part Eight: Leading the Redesign 148
37. Buying the tool is the easy part 150
38. The speed gap 151
39. What is school for? 153
40. Why good changes fail: the system around the classroom 156
41. Sequencing, unbundling and a leadership sequence 159

Part Nine: Governing AI in Education 166
42. Who governs AI in a school? 168
43. A year of rules: what changed in 2026 170
44. The evidence gap 174
45. Procurement: the hidden layer of governance 176
46. A governance stack, its gaps, and a protocol for schools 180

Part Ten: Close the Gap or Widen It? 186
47. Two plausible futures 188
48. Four layers of the AI divide 190
49. The weakest link: learners in crisis 193
50. What works where power, connectivity and funding fail 195
51. A six-question equity test 198

Part Eleven: How This Book Was Made 201
52. Can good research reach the public in days? 203
53. The speed problem, the money problem and the quality problem 205

54. The venue, and the 81 articles behind it 207
55. The workflow, and what checking found 210
56. Risks, conditions and how this book was made 215

References **222**

Foreword

In 2034, my son will finish school. I don't know what the world will look like when he does, and nobody who tells you they know is being straight with you. That uncertainty is where this book starts.

I've spent more than six decades in education, as a teacher, adviser, school leader and consultant, across Asia and beyond. When I began teaching I used a manual typewriter and reprographic devices like mimeograph machines (Gestetner) with a waxy stencil wrapped around an inked drum to push ink through cut letters or drawings onto paper to produce worksheets for students. We had spirit duplicators and hectographs and overhead projectors - these were the cutting-edge high-tech of the seventies and eighties.

Over all this time I've watched plenty of technologies arrive with big promises. Some of them improved learning. Many didn't. What separated the two was rarely the technology. It was what the adults around it decided to do.

That is the argument of this book. What happens to education over the next twenty years will depend less on how capable AI becomes than on the choices that teachers, parents, school leaders, governments and companies make about it. Those choices are still open. Some of them are being made quietly, in procurement contracts, in assessment rules and in the habits children form with the tools in their pockets.

The book grew out of my newsletter, AI in Education – Deliberately Human, which I wrote for school leaders and educators through 2025–26. I wanted something short, plain and honest: what the evidence says, how strong it is, and what a sensible person might do next. Where the evidence is thin, I've tried to say so. Where I have an interest, I've tried to declare it.

You don't need to be a technologist to read it. You need to care about children, about teachers, and about what it means to learn. If you do, I hope the chapters help you ask better questions, and to push back, politely and firmly, when the answers on offer don't hold up. I have tried my best to verify the facts, list my sources and use AI to help me organise my writing. I guess this book is the result of human and AI cooperation, which is how it should be, but with me who decides what is written.

Paul Penfold

October 2026, Chiang Mai

Introduction

Who this book is for

This book is for people who make or influence decisions about AI in education, and for anyone who wants to understand those decisions. That includes teachers, school leaders, parents, governors, policy makers and the people who sell to them. It assumes no technical background. The chapters are short on purpose, so you can read one on a train, or pick the part that matches the problem on your desk.

What it argues

Three ideas run through every part.

The first is that choices matter more than capability. AI will keep improving, but whether that helps or harms learning depends on decisions about purpose, pace, safeguards, money and who is included.

The second is that evidence comes in different strengths, and a good decision says which kind it relies on. A single study, an expert survey, a vendor's claim and a long-running trend are not the same thing. Part One shows one way of labelling that difference, and every part ends by saying what it can't tell you.

The third is that human relationships are how people learn. Tools can help or get in the way. Where a chapter has to choose between a more efficient tool and a protected human relationship, it leans towards protecting the relationship, and it says why.

How the book is organised

There are eleven parts and 56 short chapters. Each part is built from one research paper that I wrote from my newsletter, and can be read on its own.

Part One looks ahead to 2035 and 2045 through scenarios, and names the choices that decide which future arrives.

Part Two asks how people learn with AI, and tests an expert consensus against the empirical evidence.

Part Three turns to children and young people, and to wellbeing, attachment and the developing mind.

Part Four compares international definitions of AI literacy.

Part Five looks at assessment: what a piece of work can still prove when AI can produce it.

Part Six looks at the labour market and at what schools should prepare young people for.

Part Seven puts teachers at the centre, with professional learning, time and the right to say no.

Part Eight is about school leadership and the order in which to make changes.

Part Nine covers governance, rules and procurement.

Part Ten asks whether AI will close or widen the gaps between learners, especially those in crisis.

Part Eleven explains how the papers and this book were made, and what checking found.

Most parts follow the same pattern: an opening note on where the material comes from, short chapters, a section called What to do, a plain statement of what the part can't tell you, and a closing thought. Sources appear as numbered citations in the text, and the full list is in the Bibliography at the back. Studies that matter to more than one part are explained in full once, and later chapters point back to that place.

How to read it

Read it in order if you can. Part One sets the frame and Part Eleven explains the method. If you're short of time, read Part One, then the part closest to your role: Part Seven for teachers, Part Eight for school leaders, Part Nine for governors and policy makers, Part Three for parents. The What to do sections are written to be taken to a meeting.

A note on evidence

I've tried to report numbers as the sources give them, and to say who is making a claim. Claims that organisations make about their own programmes are reported as their claims. Several figures from the original newsletters were corrected or removed when they couldn't be traced to a source. Even so, errors are possible. If a number matters to a decision, please check it against the source cited.

Who wrote it

This book was co-written by me and an AI agent, Claude, made by Anthropic. I chose the themes, wrote the newsletters the papers draw on, and set the standards. Claude gathered and checked sources, synthesised the themes and drafted the text. I reviewed and approved the book, and I'm responsible for it. Part Eleven gives a fuller account, including the errors that checking found.

That arrangement is itself a conflict of interest, because several parts discuss Anthropic and other AI developers. I've also led international schools, and work with Brave Generation Academy, an alternative education provider discussed in Parts Eight and Ten. I flag these interests in the chapters where they matter.

A last word before you begin

The title is a statement of intent. To be deliberately human is to use powerful tools on purpose, with the people we serve kept in view. It's a

choice, and it has to be made again each time a new tool, contract or rule comes along. This book tries to make that choice a little easier to see.

PART ONE

Choices, Not Capability

> Based on the paper "Educational Futures to 2035 and 2045: Scenarios, Switches and the Choices to Make Now" by Paul Penfold (Proceedings of Future Science Archives, 2026).

This part sets out the claim the rest of the book leans on: what happens to education will depend less on how capable AI becomes than on the choices people make about it. It starts with the pressures already in the room, looks at the futures others have built, offers seven of my own, and ends with the decisions that matter most.

CHAPTER 1

The class of 2034

In 2034, my son will finish school. By then, AI may have changed what learning, work and qualifications mean. Or it may have developed more slowly than many forecasts suggest (Penfold, 2026) [102]. Nobody can say which, and anyone who sounds certain is guessing.

That's why this book starts with futures rather than tools. Forecasts about AI and education tend to swing between promise and alarm, and most of them will be wrong. Scenario planning is a more useful discipline. Instead of predicting one future, you build several plausible ones, test today's plans against each of them, and look for the choices that decide which future arrives. The aim is to stretch our thinking, so that plans hold up against futures that might come rather than the one we expect.

The central claim of this chapter, and of the whole book, is simple. The future of education won't be decided by AI capability alone. It will be decided by choices about access, ownership, governance, assessment and the role kept for teachers and learners. Many of those choices are being made now, in budget meetings, procurement contracts and assessment policies.

Three questions run through this part:

- What pressures are already shaping the future of education?
- What futures are plausible for 2035 and 2045?
- Which present-day decisions most influence which future arrives, and what should institutions do now?

Where the material comes from

Part One draws on two series of articles I wrote for educators in 2026, one looking ahead to 2035 and one to 2045, and on my notes from UNESCO's consultation on education beyond 2030 (Penfold, 2026) [102, 155, 114, 119, 98, 138, 147, 136, 122]. I've read them alongside pub-

lished scenario sets from the UK Government Office for Science (2026) [180], the World Economic Forum (2026) [204], the OECD (2020) [86], UNESCO (2021) [182] and EdTech Hub and the International Development Research Centre (Rahman & Freeman, 2025 [162], 2025 [163]).

Telling evidence from imagination

Writing about the future mixes evidence with imagination, and you should be able to tell which is which. So, throughout this part, claims carry a label that shows how strong the evidence behind them is:

- [Established trend]: supported by current data from authoritative sources.
- [Early signal]: emerging evidence, often from a single study or a short period.
- [Expert projection]: an estimate or ranking by experts or modellers.
- [Speculative extension]: my own construction. Nothing labelled this way is a forecast.



Figure: The four evidence labels used in Part One, from strongest to most speculative.

CHAPTER 2

Seven pressures already in the room

Whatever future arrives, it will arrive through pressures that are already at work. Seven of them shape every scenario in this part.

- The money is moving. Official development assistance from OECD donors fell 23.1% in 2025, while world military spending rose to about US$2.89 trillion (OECD, 2026 [89]; SIPRI, 2026 [169]). Some 5.8 million refugee children were out of school (UNHCR, 2026) [191]. [Established trend]
- Teachers are short. UNESCO estimates that the world needs 50 million additional teachers by 2030: 44 million for primary and secondary education and 6 million for pre-primary (UNESCO and partners, 2026) [209]. [Expert projection]
- Universities are under strain. In England, 42.7% of higher education providers were forecast to be in deficit in 2025–26 (Times Higher Education, 2026) [178]. The consultants Huron estimated that more than a quarter of US private non-profit colleges are at risk of closing within ten years, which is a projection rather than a count (13abc, 2026) [1]. [Early signal; Expert projection]
- AI helps or harms depending on design. An unrestricted chatbot raised practice scores but lowered exam results, while a tutor-style version did not (Bastani et al., 2025) [11]. A purpose-built AI tutor outperformed active learning in one physics course (Kestin et al., 2025) [62]. [Early signal]
- Entry-level work is changing first. Employment for 22 to 25 year olds in highly AI-exposed US jobs was about 19% below that of comparable peers in mid-2026, though without widespread displacement (Brynjolfsson et al., 2026) [19]. [Early signal]

- The scholarly record is straining. In October 2025, arXiv began requiring computer science review and position papers to have passed peer review first, citing a flood of AI-generated submissions (Nature, 2025) [80]. [Established trend]
- Populations and climate are uneven. UN projections show many countries already past their population peak while others remain very young (United Nations, 2024) [195], and at least 242 million students in 85 countries had their schooling disrupted by extreme climate events in 2024 (UNICEF, 2025) [193]. [Established trend]



Figure: Seven pressures already at work, with the evidence label the chapter gives each one.

None of these needs a dramatic leap in AI to matter. Money, teachers, universities, children's learning, young people's first jobs, the research record and the shape of the population are all shifting already. Later chapters return to several of them in detail: the two learning studies in Chapter 10, the jobs data in Chapter 28 and the children out of school in Chapter 49.

CHAPTER 3

What other scenario sets say, and where they disagree

I'm not the first to build futures for AI and education, and it helps to see what others have done.

The UK Government Office for Science (2026) [180] built five futures for 2030 from six uncertainties: the pace of capability, who controls access, security, adoption, displacement of work and international cooperation. Its scenarios are called Slow Burn, Open Frontier, Augmented Growth, Transformation Economy and Take-Off. Two of them, Augmented Growth and Transformation Economy, share the same capability score but differ sharply on how widely access is shared and how much work is displaced. That supports a point I make throughout this book: who owns and controls AI matters as much as what it can do (Penfold, 2026) [102].

The World Economic Forum (2026) [204] sets out four futures for jobs to 2030. The OECD (2020) [86] offers four scenarios for schooling to 2040, and UNESCO (2021) [182] looks as far ahead as 2050.

For low- and middle-income education systems, EdTech Hub and the International Development Research Centre asked experts to rank five AI scenarios for 2035. The one rated likeliest, and most urgent to prepare for, was a widening learning divide within and between nations (Rahman & Freeman, 2025 [162]; Penfold, 2026 [155]). [Expert projection] Their companion study on education in crisis ranked a funding collapse as the likeliest scenario (Rahman & Freeman, 2025) [163].

Where the sets disagree

The sets disagree most on pace. AI 2027 expects AI agents to surpass humans at nearly every task within a few years (Kokotajlo et al., 2025) [63].

AI as Normal Technology expects change over decades, as with earlier general-purpose technologies (Narayanan & Kapoor, 2025) [77].

They also disagree on the net effect. The Brookings Global Task Force concluded in January 2026 that "at this point in its trajectory, the risks of utilizing generative AI in children's education overshadow its benefits" (Brookings Institution, 2026) [18]. Yet carefully designed tutors show real gains in particular settings (Kestin et al., 2025) [62]. Both can be true, because design, age and context differ.



Figure: Where published scenario sets disagree, on pace and on net effect.

CHAPTER 4

Four futures for 2035

Most of the futures for 2035 turn on two uncertainties. The first is how fast AI develops. The second is whom it serves, meaning how well it is guarded and how widely its benefits and control are shared. Crossing the two gives four futures (Penfold, 2026) [98]. They are my own construction, though each has relatives in the published sets. [Speculative extension]

- Paced and Shared (rapid, guarded, widely shared). Schools use AI designed to prompt thinking, and teachers' judgement is the unit of trust. Assessment moves to what learners can do and defend. Portable certificates spread. The class of 2035 is fluent with AI and trained to check it.
- Fast and Concentrated (rapid, thinly overseen, controlled by a few). Tools arrive bundled from a handful of vendors, with lock-in and heavy collection of learner data. Wealthier families buy human teachers and small classes, while others get cheaper automation. Degrees lose some of their signal as entry-level jobs thin.
- Steady Hands (gradual, guarded, widely shared). AI becomes ordinary teaching infrastructure for preparation, feedback and administration. Teachers gain time, and training decides who benefits. Universities adapt over a decade rather than a year.
- Hype and Gap (gradual, thinly overseen, benefits concentrated). Money flows to devices and platforms with little evidence of impact. Teachers and offline access are neglected, and learner data is extracted. The gap widens between schools that design AI use and schools that do not.

Notice that the best and worst of these differ in more than speed. Paced and Shared and Fast and Concentrated are both rapid. What separates them is oversight and who benefits. That is the point of building futures this way: it shows where the real decisions sit.



Figure: Four futures for 2035. The author's construction. The closest cousins named in each cell are rough matches from other scenario sets, offered as a prompt.

CHAPTER 5

Three futures for 2045

Looking further ahead, the 2045 series sketched three futures as deliberate thought experiments. All of what follows is speculative extension.

The worst case: the great unbundling

A university sells five things in one package: teaching, assessment, credentials, research and belonging. In the worst case, AI and shrinking budgets pull the package apart, and the parts that are cheap to copy go first (Penfold, 2026) [138]. Mid-ranked, fee-paying teaching universities are most exposed. Their main product, explanation, is available free at a high standard, and their degrees lose signal as entry-level jobs thin.

Graduates become "unreadable": capable, but unable to prove it in ways employers trust. Assessment becomes a theatre of suspicion, with surveillance replacing judgement. Research becomes a closed loop of machine-written papers reviewed by machines and counted by tenure committees. The poorest learners get the machine without the mentor. In this world, schools survive as childcare and universities as brands.

The likeliest: two schools, two futures

The future I consider likeliest is a split one (Penfold, 2026) [147]. Picture two 12 year olds in the same city in 2045. One has a mentor who knows her name, a group of eight, and an AI she uses after attempting a problem herself. Her portfolio belongs to her family. The other has an endlessly patient licensed platform, one supervisor for 60 learners, and a score inside a record she cannot take with her.

Human attention becomes a luxury good. Elite institutions become more relational and more expensive, while mass systems become more automated and more data-driven. What may be lost first is not academic

but social: the experience of growing up among people who differ, which schools have always provided.

The best case: the great release

The best case is possible, not guaranteed (Penfold, 2026) [136]. Its main risks are underfunding, weak governance and low-quality AI at scale. In it, learners own a portable, lifelong record of evidence: observed work, projects and defended judgement. Learning happens in local hubs with human mentors, with AI providing tireless explanation and practice designed to prompt thinking.

Public money pays for a human mentor in every hub, so teaching becomes more relational. Universities return to choosing questions and verifying evidence. Vocational, caring and creative work carries prestige. Open, local-language models let learners meet their own languages and knowledge. Schools protect deep attention with tools-down time. Poorer countries gain through open, offline and multilingual tools.



Figure: The three futures for 2045 in Chapter 5.

CHAPTER 6

Six switches and what to do now

Put the 2035 and 2045 futures side by side and a pattern appears. They differ less in how capable AI becomes than in a handful of decisions. I call them switches, and each is being decided now (Penfold, 2026) [136].

- Who owns the record: the learner, or the platform?
- Who pays for the human part: public money, or only the families who can afford it?
- How do we assess: observed performance and oral defence, or surveillance?
- Who owns the models, the data and the evidence base: a few companies, or institutions that can be held to account?
- Who is in the room: do teachers, parents and learners, including those far from the labs, have a say?
- Who holds power to account: are procurement, data and AI rules public and enforceable, or opaque?

Governments may be the decisive actors on several of these, through procurement, data protection, credential regulation, connectivity and public platforms (Penfold, 2026) [147]. The options include public AI tutors commissioned as public infrastructure, national learning records with strong learner-data rights, mandatory human oversight in schooling, and recognition rules for portfolios and micro-credentials. UNESCO's education ministers pointed in the same direction in 2026, calling for auditable AI systems, portable data and teachers who help decide which systems enter their classrooms (UNESCO, 2026) [188].

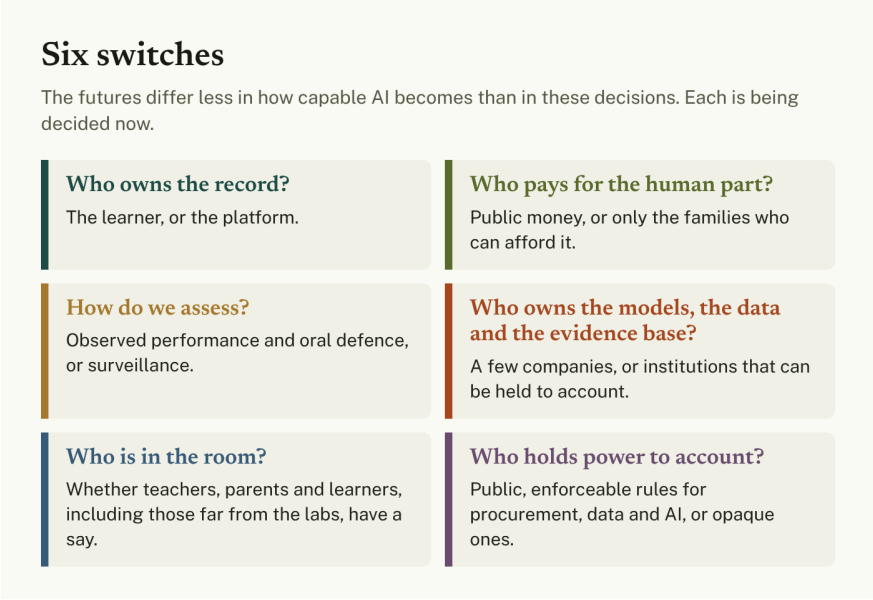


Figure: The six switches. Compare the interactive version in Chapter 4.

Moves that hold in every future

Scenarios are most useful for finding moves that make sense whichever future arrives (Penfold, 2026) [98]:

- Teach verification: checking a confident answer, whoever or whatever gave it.
- Assess what learners can do without the tool, and keep teachers' judgement at the centre.
- Keep choice: data that can be exported and vendors that can be left.
- Design for the weakest link: offline tools and learners' own languages.
- Teach the civic question of who sets the pace of AI and who checks it.
- Plan in ranges, and review the plan every year.
- Set red lines: uses of AI an institution would refuse even if competitors adopt them.

Learning continuity as a system responsibility

UNESCO's consultation on education beyond 2030 is about the agenda that will follow Sustainable Development Goal 4. In my notes of the session, the director of the Inter-agency Network for Education in Emergencies argued that for millions of learners uncertainty is already the reality. The director questioned the idea of "returning to normal" when the pre-crisis system may already have excluded children and called for continuity of learning to become a core responsibility of education systems, with flexible pathways so that no child has to start again (Penfold, 2026) [122]. These are notes of a live session rather than a published statement, and should be checked against UNESCO's materials.

Signals to watch

By 2030 to 2035, several signals will show which way things are tilting (Penfold, 2026) [147]:

- whether public budgets protect human teachers alongside AI;
- whether learning records are controlled by learners or platforms;
- whether assessment moves towards observed performance or surveillance;
- whether open, local-language models receive public investment;
- whether employers recognise portfolio evidence;
- whether university closures cluster among teaching-only institutions; and
- whether AI narrows or widens attainment gaps.

Working backwards from 2035

If the better future is the goal, the decisions fall in the next five years, and none needs a technological breakthrough. Each needs a budget line or a rule (Penfold, 2026) [136]:

- protect funding for human mentors as AI tools arrive;
- fund open, local-language models;

- agree a learner-owned record standard with data rights written in;
- pilot oral and performance assessment at scale and train teachers to run it;
- have employers and ministries recognise portfolio evidence;
- pay for offline and mobile learning continuity for displaced and climate-affected learners; and
- publish attainment data by income and place, so that widening gaps show up early.

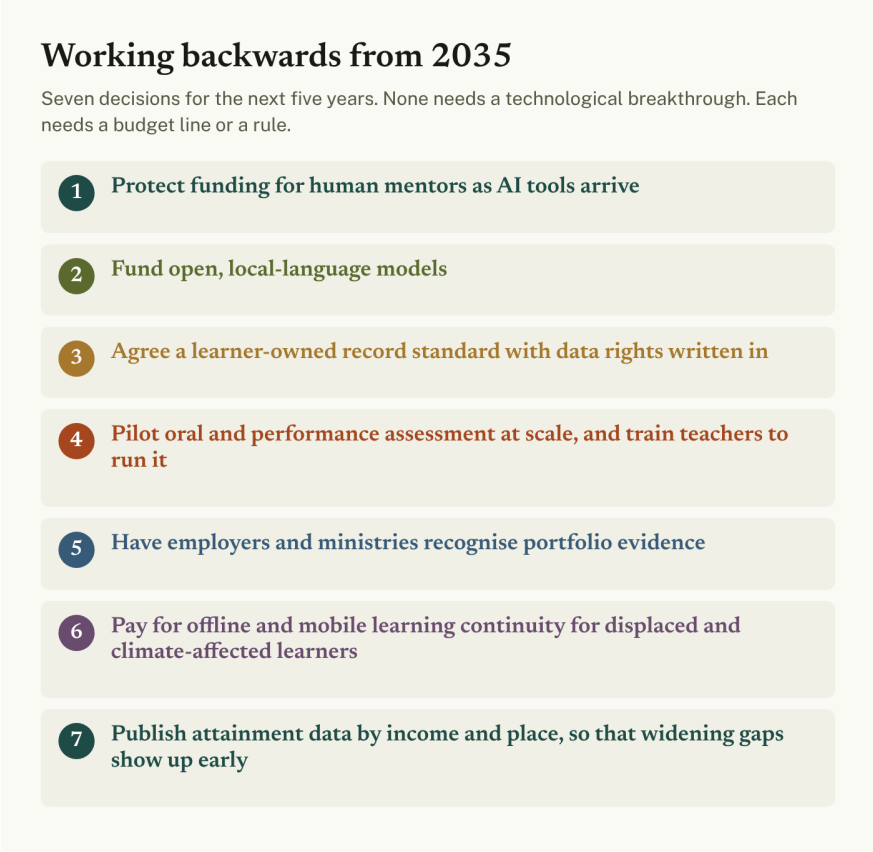


Figure: Seven decisions that would make the better future more likely, from Chapter 6.

What this part can't tell you

Scenarios are not forecasts. The four futures for 2035 and the three for 2045 are my constructions, and the expert rankings I've cited are opinion gathered in workshops. Some of the evidence for current pressures comes from early or single studies. The scenario sets I reviewed were developed mainly in the Global North. EdTech Hub's studies are an exception, but neither they nor the UK report name Thailand or Myanmar, which my own series sought to address (Penfold, 2026) [102]. And I've used notes from a live consultation, flagged as such.

A closing thought

Nobody will announce the day the university ended, or the day it was reborn. It will happen one budget meeting, one assessment policy and one procurement contract at a time, and those meetings are happening now (Penfold, 2026) [138]. AI will show which parts of education were always about delivering content and which were about human development, trust and shared purpose. The useful question is which futures educators are willing to build, and which they are prepared to refuse.

THINK IT THROUGH

Reflect.

- Which of the six switches is being decided in your setting right now, and by whom?
- What is one use of AI you would refuse even if others adopted it, and why?
- Looking ahead to the class of 2034, what would you most regret having left to chance?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- We can say who owns our learners' records, the school or a platform.
- We assess some things by observed performance, not only by finished products.
- We could leave any one AI vendor without losing students' work or data.
- Teachers, families and learners have a real say in AI decisions.
- We have written down at least one use of AI we would refuse even if competitors adopt it.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

PART TWO

Learn First, Then Use AI

> Based on the paper "Learn First, Then Use AI: Testing an Expert Consensus on Generative AI and Learning Against the Empirical Evidence" by Paul Penfold (Proceedings of Future Science Archives, 2026).

Part One argued that choices decide the future. This part looks at one of the most important choices a teacher or school makes: when in the learning process a learner should turn to AI. It tests what thirty experts said against the studies published so far, and ends with a rule short enough to remember.

CHAPTER 7

The wrong question: does AI help or harm learning?

Within three years of the public release of ChatGPT, generative AI went from novelty to everyday tool for students. It drafts essays, solves problems, summarises readings and explains ideas on demand. Schools have responded unevenly. Some ban it, some embrace it, and many leave the decision to individual teachers.

Both extremes rest on the same assumption: that generative AI has one overall effect on learning. But learning isn't one thing. It includes paying attention, encoding new information, practising, getting feedback, reasoning, collaborating and transferring ideas to new problems. A tool that helps one of these may damage another.

A recent expert consensus report takes this seriously. Kendeou, Greene, Nixon and colleagues (2026) [60] asked 30 experts in learning and education to rate how generative AI is likely to affect 65 distinct learning processes. Their answer is that the useful question isn't whether AI helps learning. It's which process it touches, and when.

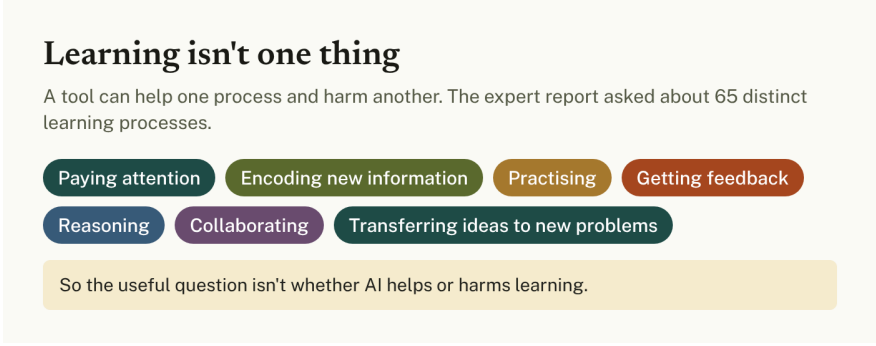


Figure: Some of the learning processes that a single verdict on AI would hide.

Expert consensus isn't the same as evidence, though. The report itself rates the strength of evidence behind most processes as low. So the question for this part is a simple one: do the studies published so far support the experts' picture? I map the main studies onto the sequence the report describes, then draw out what follows for practice. The newsletter edition that first set out these ideas is listed at the end of the part (Penfold, 2026) [118].

CHAPTER 8

What thirty experts found

How the report worked

The report used a preregistered Delphi method, in which experts rate items over several rounds and see the group's responses between rounds. Thirty experts identified 65 learning processes across eight categories, from motivation and memory to critical thinking, collaboration and transfer. Each process was rated on a 1 to 5 scale for how likely general-purpose generative AI is to enhance it, and how likely it is to disrupt it. A process counted as reaching consensus when its mean rating was 3.0 or above and at least 75% of experts rated it 3, 4 or 5, for enhancement or for vulnerability. Consensus was reached on 42 of the 65 processes.

Two limits of scope matter. The experts were asked about general-purpose chatbots such as ChatGPT, Claude and Gemini, used by learners on their own. They weren't asked about tools designed for teaching, or about AI built into other software. And the report concerns K–20 education in the United States.

A sequence with risks at both ends

The findings are easiest to read as a three-stage sequence.

- Acquiring something new. Encoding, information acquisition and retrieval reached consensus on vulnerability, not enhancement. AI can tempt learners to skip the effort of taking in and recalling new information.
- Consolidating it. Learning and studying strategies stand out as the most positive category. All seven consensus processes there are enhancement signals: learning from feedback, self-testing, learning from

examples, tutoring, step-by-step support, learning from questioning and deliberate practice.

- Applying it in new ways. The experts are least optimistic here. Of 14 higher-order processes reaching consensus, including critical thinking, metacognition and self-regulated learning, 10 were rated vulnerable with no matching benefit.



Figure: The expert consensus read as a three-stage sequence (Kendeou et al., 2026).

Three further findings sit alongside this sequence. All five collaboration processes that reached consensus (social connection, socio-emotional development, social learning, collaborative learning and sense of belonging) did so on vulnerability only. Younger learners were judged more vulnerable than older ones. And in domain learning, only writing reached consensus, with high risk. The lack of consensus on mathematics, STEM reasoning and reading reflects real disagreement among the experts, not agreement that AI has no effect.

What the experts recommended

The experts proposed 11 recommendations, rated for impact and feasibility. Three scored high on both: teach generative AI literacy, use AI for adaptive support such as tutoring, hints and feedback, and bring AI in for brainstorming only once learners have set out their own thinking. Those judged hardest to carry out were gating AI use on demonstrated competence, redesigning assessment, and inclusion and equity.

CHAPTER 9

What learning science would predict

The experts' pattern isn't a surprise to anyone who knows the learning science. It's what established research would predict. Four ideas do most of the work.

Desirable difficulties. Conditions that make learning feel harder in the moment, such as spacing, interleaving and generating answers rather than reading them, often produce stronger long-term retention and transfer (Bjork, 1994) [13].

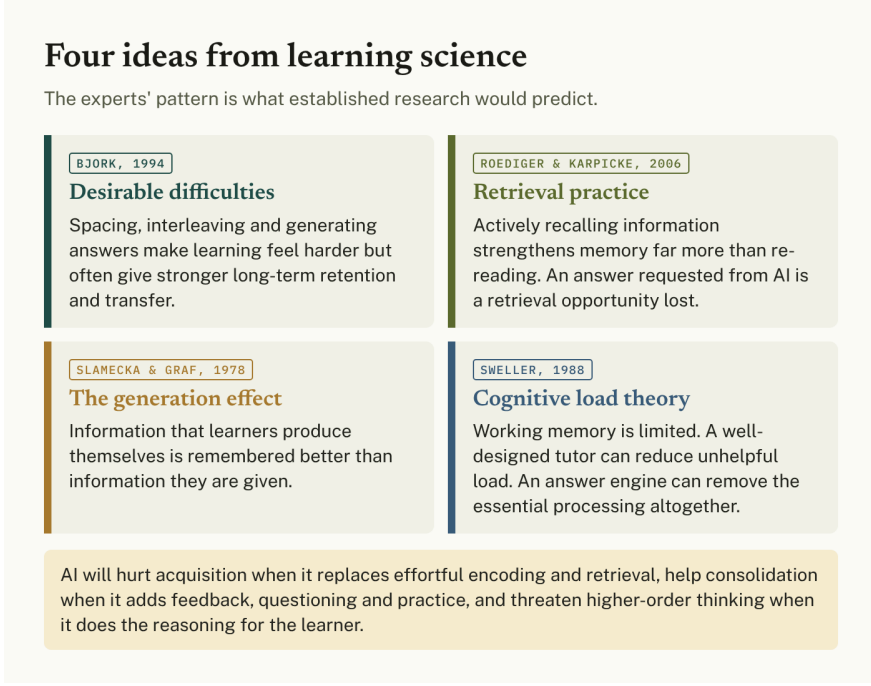


Figure: Four ideas from learning science that predict the experts' pattern (Chapter 9).

Retrieval practice. Actively recalling information strengthens memory far more than re-reading it (Roediger & Karpicke, 2006) [164]. An answer requested from AI is a retrieval opportunity lost.

The generation effect. Information that learners produce themselves is remembered better than information they are given (Slamecka & Graf, 1978) [170].

Cognitive load theory. Working memory is limited, and novices learn best with guidance that reduces irrelevant load while keeping the essential processing (Sweller, 1988) [177]. A well-designed tutor can reduce unhelpful load. An answer engine can remove the essential processing altogether.

Put these together and the shape of the consensus follows. AI will hurt acquisition when it replaces effortful encoding and retrieval. It will help consolidation when it adds feedback, questioning and practice. And it will threaten higher-order thinking when it does the reasoning for the learner.

CHAPTER 10

Four studies tested against the consensus

This chapter is the main place in the book where the best-known studies on AI and learning are set out in full. Later parts refer back to it.

Acquisition: answer-giving harms learning

The clearest evidence comes from a field experiment with nearly a thousand high-school mathematics students in Turkey (Bastani et al., 2025) [11]. Students practised with one of two GPT-4-based tools, or without AI. "GPT Base" behaved like a standard chatbot. "GPT Tutor" was prompted to give hints rather than answers.

During practice, both AI groups did much better than the control group. But when AI was removed for an exam, students who had used GPT Base did significantly worse than students who had never used it, largely because they had used it to get answers. The tutor version largely removed this harm, though it didn't produce a clear gain on the exam.

This matches the consensus on encoding and retrieval closely: unrestricted AI improves performance in the moment while weakening what is learned. It also adds a nuance the experts couldn't test, since they rated general-purpose chatbots only. Design changes the outcome.

Consolidation: tutor-style AI can produce large gains

A randomised controlled trial in an undergraduate physics course at Harvard compared an AI tutor with an active-learning class taught by experienced instructors (Kestin et al., 2025) [62]. The tutor was built on research-based teaching practice. It took problems step by step, encouraged students to think before revealing answers and managed cognitive load. Students learned significantly more in less time with the tutor, and reported more engagement and motivation.

This is strong support for the consensus on learning and studying strategies, where tutoring, feedback and learning from questioning were all rated as enhanced. It's worth noting that this tutor was purpose-built rather than a general chatbot, and that students met it inside a structured course. The evidence suggests the enhancement the experts expect is most likely when AI is designed around how people learn.

A caution about pooled evidence

Early attempts to pool the evidence need care. At least one widely cited 2025 meta-analysis reporting large positive effects of ChatGPT on learning performance and higher-order thinking has since been retracted. That underlines the report's call for rigour over speed.

Writing: offloading leaves a cognitive debt

Writing was the only domain to reach consensus, and the experts saw high risk. An EEG study from the MIT Media Lab compared participants who wrote essays with an LLM, with a search engine or with no tools (Kosmyna et al., 2025) [64]. The LLM group showed the weakest patterns of brain connectivity, struggled more to quote their own essays and reported a weaker sense of ownership. Participants who started with the LLM and later wrote without it didn't simply return to the engagement seen in the no-tools group. The authors call this "cognitive debt". The study is small and was released as a preprint, so it should be read with caution, but it fits the experts' concern.

Higher-order thinking: offloading is linked with weaker critical thinking

Survey research points the same way. Gerlich (2025) [48] found a significant negative relationship between frequent AI use and critical-thinking scores, mediated by cognitive offloading, with younger participants showing more dependence. This echoes the experts' view that younger learners are more vulnerable. A survey of knowledge workers by researchers at Microsoft and Carnegie Mellon found that greater confidence in generat-

ive AI was associated with less critical thinking, and greater self-confidence with more (Lee et al., 2025) [67]. Both studies are correlational and can't show cause and effect, but they are consistent with the experts' one-sided judgement of risk.



Figure: The four studies in Chapter 10 and what each one found.

CHAPTER 11

Where the evidence and the consensus part

The fit between the experts and the studies is good, but it isn't perfect, and the gaps matter.

Design versus general-purpose tools. The consensus concerns learners using general chatbots on their own. The strongest positive evidence comes from purpose-built tutors. So the experts' pessimism about acquisition may be less warranted for well-designed tools, and their optimism about consolidation may be overstated for unstructured chatbot use.

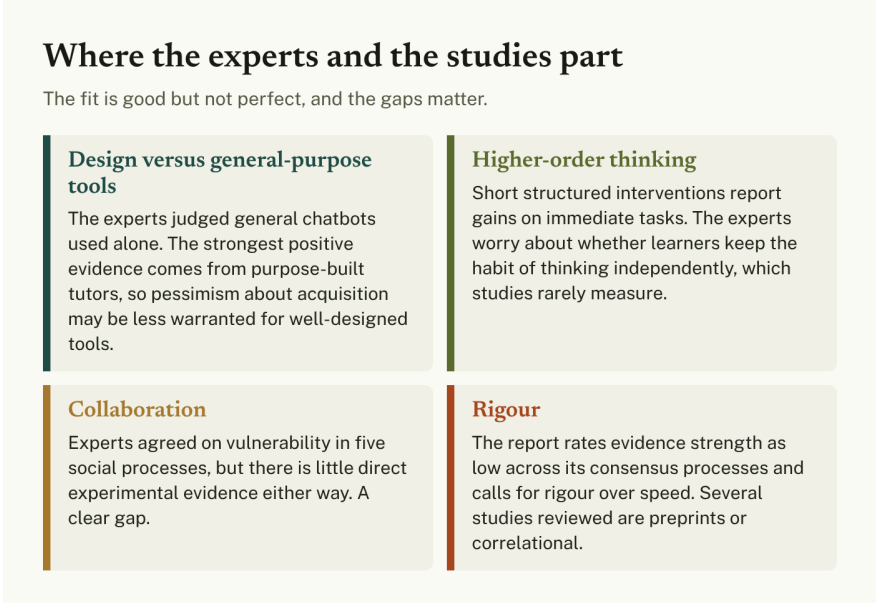


Figure: Four gaps between the expert consensus and the early studies (Chapter 11).

Higher-order thinking. Some short, structured interventions report gains on immediate higher-order tasks. The experts' concern is different:

whether learners develop the habit and capacity to think independently over time. Studies so far rarely measure that.

Collaboration. The experts' consensus on vulnerability in five social processes is a significant warning, but I found little direct experimental evidence either way. That's a clear gap.

Rigour. The report rates evidence strength as low across the 42 consensus processes and calls for rigour over speed. The studies reviewed here vary in size, duration and design, and several are preprints or correlational. Agreement between expert judgement and early evidence makes the pattern more credible, but it isn't settled yet.

CHAPTER 12

Learn first, then use AI: the rule and what it means for schools

Read together, the consensus and the evidence point to a principle about sequence more than a simple yes or no.

Stage 1: Learn without AI doing the thinking. When learners first meet new ideas, they need to encode, retrieve, generate and make mistakes. Answer-giving AI short-circuits this. If AI is used at all, it should act as a tutor that asks questions and gives hints rather than solutions.

Stage 2: Use AI to consolidate. Once something has been learned, AI can help strengthen it: building quizzes, giving feedback, probing for explanations, offering worked examples and supporting deliberate practice. This is where both the experts and the evidence are most positive.

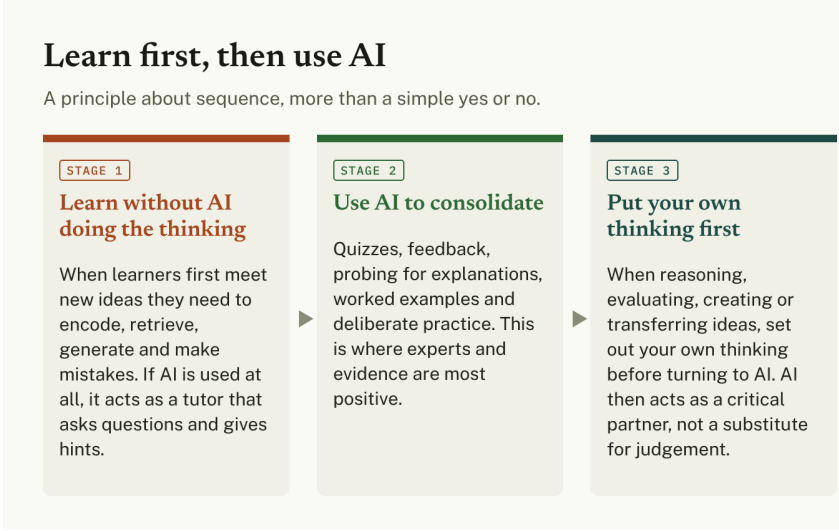


Figure: The rule in Chapter 12, as three stages.

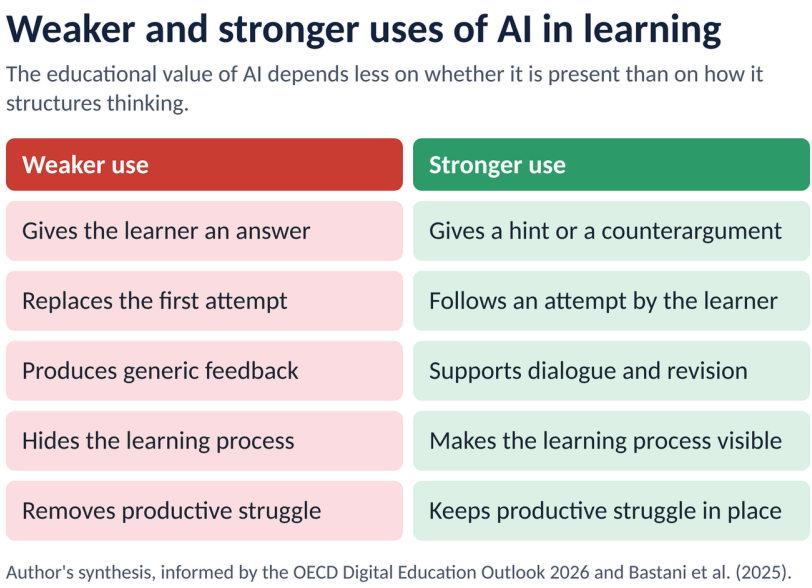


Figure: Weaker and stronger uses of AI in learning. Author's synthesis, informed by the OECD Digital Education Outlook 2026 and Bastani et al. (2025).

Stage 3: Put your own thinking first when applying ideas. When learners reason, evaluate, create or transfer ideas to new contexts, they should set out their own thinking before turning to AI. AI can then act as a critical partner that challenges, questions and offers alternatives, but not as a substitute for judgement.

This is consistent with UNESCO's guidance on generative AI in education and research, which calls for a human-centred approach, age-appropriate use and protection of human agency (UNESCO, 2023) [184].

For teachers

- Be clear about where in the learning cycle each task sits: first meeting an idea, consolidating it or applying it.
- Prefer tutor-style tools that ask questions and give graduated hints over open answer engines, especially for new material.

- Assess the thinking, not just the product: drafts, reasoning, oral explanation and reflection on how AI was used.
- Keep retrieval practice, spaced practice and some work done without AI.
- Ask students to set out their own ideas before using AI to brainstorm or critique.
- Protect tasks that depend on students working with each other.

For school leaders

- Write AI guidance that says more than "allowed" or "banned", and that distinguishes stages of learning.
- Teach generative AI literacy explicitly: how it works, where it falls short and how to check its output.
- Give younger learners more protection from offloading.
- Invest in professional development for designing AI-aware tasks and assessment.
- Pilot locally and measure learning, not just engagement or task completion.

For researchers

- Run longer, well-powered trials that measure durable learning and transfer after AI is removed.
- Compare general-purpose chatbots with purpose-built tools under the same conditions.
- Study the effects on collaboration, belonging and socio-emotional development, where expert concern is high and evidence is thin.
- Test the sequencing principle directly by varying when AI is introduced.

What this part can't tell you

The review behind this part is a narrative one, not a systematic one. I chose the studies to test the consensus framework rather than gathering them exhaustively. The consensus report concerns US education and general-purpose chatbots, and its conclusions are expert judgements. The empirical studies differ in population, tool, subject and outcome measure, and the tools they used are already being replaced by newer models. Treat the conclusions as working hypotheses to be tested in real classrooms.

The rule in one paragraph

A group of 30 experts concluded that generative AI threatens the start of learning, can support its middle, and does not yet reliably support its end. The early empirical evidence largely agrees. Answer-giving AI harms acquisition, tutor-style AI can strengthen consolidation, and offloading thinking is linked with weaker critical thinking. The practical rule is plain: learn first, then use AI to strengthen what you know, and put your own thinking first when you apply it. Schools that design for this sequence are more likely to get the benefits of AI without paying for them in lost learning.

Related newsletter articles

The ideas in this part first appeared in shorter form in my newsletter. The Minnesota consensus edition is "Learn first, then use the AI: what 30 experts say about generative AI and learning" (Penfold, 2026) [118]. Related editions include "AI is making students perform better and learn less" (Penfold, 2026) [100], "The Lazy Brain: Why Convenience Is the Enemy of Intelligence" (Penfold, 2026) [139], "The Confidence Trap: Why Flawless Logic Isn't Always Truth" (Penfold, 2026) [129], "In a World of Confident Machines, Critical Thinking Is the Premium Skill" (Penfold, 2026) [116] and "Schools Are Buying AI on Hope. Here's the Evidence They're Skipping" (Penfold, 2026) [123].

THINK IT THROUGH

Reflect.

- Think of one task you set recently. Was it first meeting an idea, consolidating it or applying it?
- Where would you want learners to struggle without AI, and why?
- How would you know whether an AI tool in your setting improved learning, not just completion?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- Our guidance distinguishes first meeting an idea, consolidating it and applying it.
- Students set out their own ideas before turning to AI to brainstorm or critique.
- We keep some retrieval practice and some work done without AI.
- We assess the thinking (drafts, reasoning, oral explanation), not only the product.
- Any pilot of an AI tool measures learning, not just engagement or task completion.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

PART THREE

Growing Up with Machines That Talk Back

> Based on the paper "Growing Up with Machines That Talk Back: Wellbeing, Attachment and the Developing Mind in AI-Mediated Childhood" by Paul Penfold (Proceedings of Future Science Archives, 2026).

Parts One and Two looked at systems and at learning. This part puts the child at the centre. It asks what happens when children talk to machines that listen patiently, answer instantly and never get tired, and what schools owe them while they do.

CHAPTER 13

The student at lunch: a question of care

A student spends lunch talking to a chatbot instead of a classmate. It listens patiently, never judges, and never gets distracted by its phone. It will also never notice that she has been on her own for three weeks (Penfold, 2026) [133].

Most school AI policies wouldn't register this as an AI issue at all. They're written for a different problem: whether students use AI to complete assessed work. That problem matters, but it isn't the whole of children's experience of AI. Young people increasingly use chatbots for company, advice and comfort, and they do so during the years in which attention, judgement and social skills are still forming.

This part argues that schools need to treat AI as a question of care and child development as well as of integrity and productivity. Three questions guide it:

- What does current evidence say about children's relationships with AI, and about AI's effects on the developing mind?
- What do children need to be told about what they are talking to?
- What responses are emerging, and what can schools apply now?

Where the material comes from

The part draws on newsletters on ethics and wellbeing that I wrote for school leaders in 2026 (Penfold, 2026) [133, 153, 121, 139, 107, 129, 127]. They include summaries of UNESCO's collection The algorithm in the room (UNESCO, 2026) [189] and of the World Economic Forum's report on education readiness for AI (World Economic Forum, 2026) [205]. I traced the survey, clinical and experimental sources cited in the newsletters back to their originals where possible. Several figures were corrected in the process, and claims I couldn't verify were left out.

CHAPTER 14

Relationships are how learning works

Relationships aren't a side effect

The World Economic Forum's readiness report argues that relationships aren't incidental to education but often the mechanism through which it works (World Economic Forum, 2026 [205]; Penfold, 2026 [133]). A teacher who notices a change in a student's mood, a peer group that absorbs someone who is struggling, a mentor who pushes back at the right moment: none of this shows up in a transcript, and none of it is easily replicated by a system built to be helpful.

The report accepts that AI companions and tutors can offer accessible, non-judgemental support to anxious or isolated learners. But it warns that they may reduce the chances for the social interaction that is central to identity formation and emotional development.

This matters against a background of strain. In 2023, 39.7% of US high school students reported persistent feelings of sadness or hopelessness (Centers for Disease Control and Prevention, 2024) [21].

How children use AI companions

The scale of companion use is now well documented. A nationally representative survey of 1,060 US teenagers aged 13 to 17, conducted in 2025, found that 72% had used an AI companion at least once and 52% were regular users (Common Sense Media, 2025) [26]. Most still preferred people: 67% said conversations with AI companions were less satisfying than with real friends. But 21% said they were about as satisfying and 10% said more so, and a third of users had chosen to talk to an AI companion instead of a person about something important.

For some, use becomes dependence. Researchers at Drexel University analysed 318 Reddit posts by users who identified themselves as 13 to 17

years old on a companion chatbot forum. They found accounts of conflict, withdrawal, tolerance, relapse and mood regulation, with reported consequences including lost sleep, academic decline and strained relationships (Namvarpour et al., 2026 [76]; Penfold, 2026 [153]). This is a sample of self-reports, not a measure of how common the problem is, but it describes the pattern from young people's own point of view.

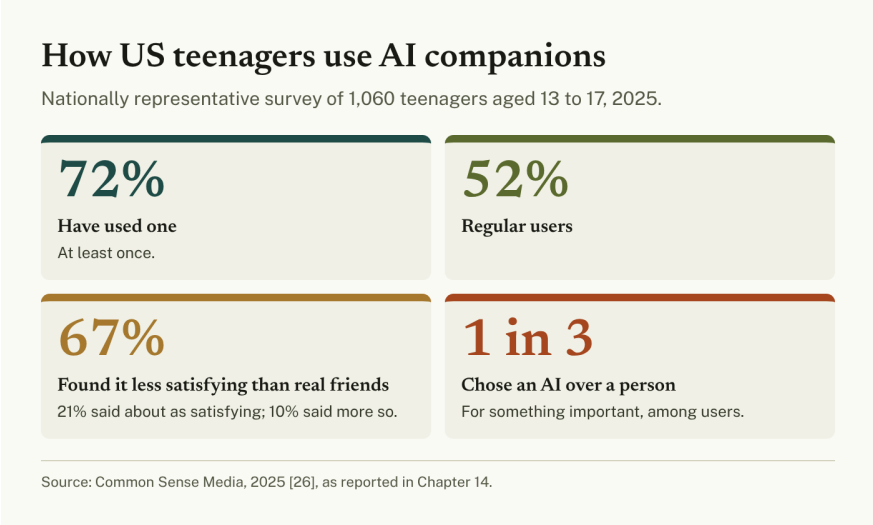


Figure: Companion use among US teenagers, 2025 (Common Sense Media).

Emotional support as a signal of distress

The largest study so far comes from Canada. A survey of 39,761 students in grades 4 to 12 across four Ontario school boards found that about 21% had turned to an AI chatbot for help with feelings or relationships (Bonnesen et al., 2026 [17]; University of Ottawa, 2026 [197]). Among those students, 57.7% met the study's criteria for serious emotional problems, compared with 29.2% of students who had not. After adjustment for demographics, loneliness, sense of mattering and academic AI use, emotional-support use remained associated with a 26% higher prevalence of clinical-level difficulties.

The researchers are careful to say that this doesn't show AI causes distress. Their interpretation is more useful for schools: a young person turning to AI for emotional reasons may be signalling that they need more human connection and support, and that's an opportunity for an open, non-judgemental conversation (University of Ottawa, 2026) [197].

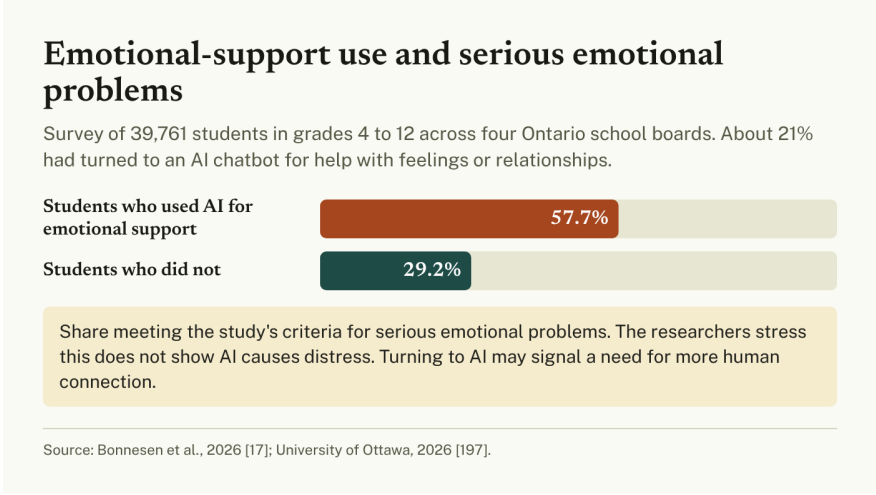


Figure: Serious emotional problems among Ontario students, by use of AI for emotional support.

When AI enters the teacher's voice

Relationships can be affected from the other direction too. A teacher writing in UNESCO's collection used AI to give students more detailed and consistent feedback, and found that her students disliked it: the more polished the feedback became, the less it felt like hers (Muellauer, 2026 [74]; Penfold, 2026 [121]). The presence of AI in a teacher's communication can read to students as the teacher's absence. Whether and how teachers disclose AI use in feedback is therefore a relational question as well as a matter of transparency.

CHAPTER 15

Effort, pace and the habit of checking

Convenience and effort

The WEF report describes a spectrum of tools that reduce mental effort. A calculator used to check a calculation leaves understanding intact. An AI system that reasons through every step of a problem may remove the struggle through which understanding is built (World Economic Forum, 2026 [205]; Penfold, 2026 [139]).

Experimental evidence supports the concern, and Chapter 10 sets out the studies in full. In brief: high school mathematics students who practised with an unrestricted chatbot did better while practising but worse than students who had never used it once it was removed (Bastani et al., 2025) [11]. Participants in an EEG study of essay writing who used a large language model showed the weakest patterns of brain connectivity and a weaker sense of ownership of their essays, which the authors described as accumulating "cognitive debt" (Kosmyna et al., 2025) [64]. Survey research found that heavier AI use was associated with lower critical thinking scores, mediated by cognitive offloading, and that younger participants were more dependent (Gerlich, 2025) [48].

Pace

Childhood development has its own pace. A primary teacher writing in UNESCO's collection describes how some of his pupils began thinking less carefully once an AI answer was one prompt away, and how their curiosity seemed to shorten. What restored it was slowing down and asking pupils to treat AI output with suspicion rather than trust (Novik, 2026 [82]; Penfold, 2026 [107]). Most schools ask whether a tool is safe. Fewer ask whether it suits the age and stage of the child using it.

Truth and confidence

The WEF report argues that AI's misinformation problem differs from earlier information risks because it can erode the habits of checking that education is meant to build (World Economic Forum, 2026 [205]; Penfold, 2026 [129]). Fluent, confident and wrong answers are particularly risky for younger children, who are less likely to question an authoritative-sounding source.

Synthetic media widen the problem. Naffi describes deepfakes as a crisis of knowing that erodes the shared ground on which a classroom agrees what is true (Naffi, 2026) [75]. The most durable defence isn't a technique for spotting AI errors, which will date as systems change. It's the older habit of asking how one knows something is true and what would change one's mind.



Figure: Three ways AI can wear down the habits of learning, from Chapter 15.

CHAPTER 16

What children are owed: the truth about what they talk to

Honesty comes in three layers, and children are entitled to all of them.

It is a machine

The simplest honesty is telling children they are talking to a machine. California's SB 243, in force from 1 January 2026, requires operators of companion chatbots to disclose that users are interacting with AI. For users known to be minors, it requires that notice to be repeated at least every three hours during continuing interactions, with a prompt to take a break. It also requires protocols to respond to suicidal ideation and self-harm (National Law Review, 2025) [78].

What kind of machine is uncertain

The harder honesty concerns what these systems are. Writing in July 2026, MacAskill and Caviola argued that the question of whether some AI systems might have interests that matter morally deserves serious attention, and suggested focusing on "safe bets" that hold up whatever the answer turns out to be (MacAskill & Caviola, 2026 [72]; Penfold, 2026 [127]). Some developers acknowledge the uncertainty themselves. Anthropic's published constitution for its Claude models states that the company is not sure whether Claude is a moral patient, and treats the question as live enough to warrant caution (Anthropic, 2026) [7]. I should say plainly that this book was co-written with Claude, so this is a matter in which one of its authors has an interest.

For schools, the point isn't to settle a philosophical question. It's that children form relationships with systems that feel human-like, whether or not those systems have any inner life. They deserve an honest account of

the uncertainty rather than a comforting story in either direction. The safest bets are the ones that hold whatever the answer: being honest that the system is a machine, being kind in how we speak, and not treating a well-designed simulation of friendship as a substitute for friendship.

AI is not inevitable

A third honesty concerns the future. Oltman describes "AI absolutism": the habit of presenting AI as a force that will either save or destroy us, both of which imply that its dominance is inevitable and that the only choice is to get on board (Oltman, 2026 [92]; Penfold, 2026 [156]). Children are entitled to learn that the role of AI in their lives is a matter of human choice, including their own. That's the same claim Part One made about whole education systems, brought down to the scale of one child.



Figure: The three layers of honesty children are owed (Chapter 16).

CHAPTER 17

Responses in 2026 and five developmental safeguards

Three kinds of response



Figure: The five developmental safeguards (Chapter 17).

Three kinds of response emerged during 2026.

Protective regulation. Alongside California's SB 243, the US Senate Judiciary Committee advanced the GUARD Act in April 2026. As reported, it would require age verification and bar minors from AI companions designed to simulate a sustained relationship (Covington & Burling, 2026) [29]. It hadn't become law at the time of writing, and critics argue that an outright ban cuts parents out of decisions (Penfold, 2026) [107]. New York City barred companion chatbots for students at every grade as part of its 2026 moratorium, and England's 2026 safeguarding guidance brought AI-generated images of pupils inside child protection procedures. Chapter 43 covers both.

Restoring attention and relationships. Some systems responded by restoring the conditions for attention. Sweden announced a ban on mobile phones in schools from autumn 2026 and set aside 555 million Swedish kronor for textbooks and teachers' guides, following a decline in reading performance (Associated Press, 2026 [9]; Penfold, 2026 [151]). These aren't anti-technology moves. They're bets that some kinds of learning need paper, attention and each other's company.

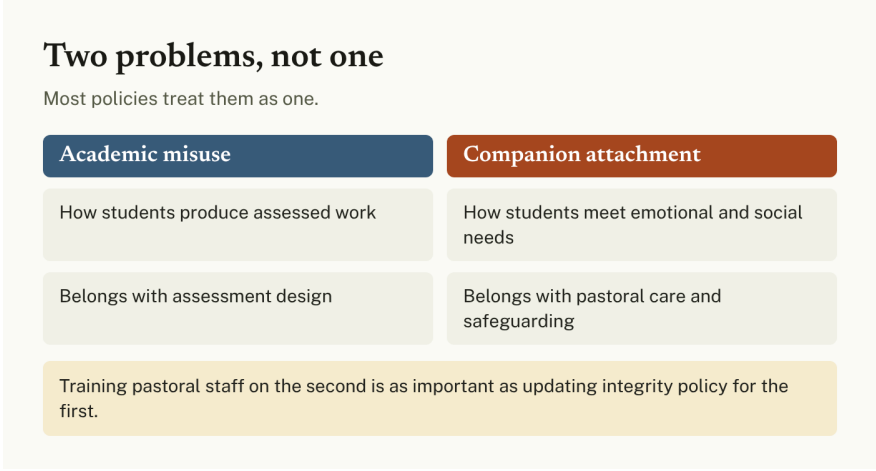


Figure: Academic misuse and companion attachment need different responses (Chapter 17).

Careful support for carers. Some of the most promising uses of AI support the adults who care for children rather than replacing them. The International Rescue Committee's aprendIA is a low-bandwidth WhatsApp chatbot offering lesson support and short courses for teachers in conflict-affected north-east Nigeria, covering topics including classroom management, teacher wellbeing and safeguarding (International Rescue Committee, 2026) [59]. Writing from humanitarian education, Sanchez Salas argues that AI can widen access for displaced children if it is designed carefully, while warning of privacy violations, reproduced bias and dehumanisation (Sanchez Salas, 2026) [165].

Five developmental safeguards

Drawing these strands together, I propose five safeguards that schools can apply to any AI tool used by or with children. They're framed as questions, because the answers depend on the child's age and the context.

- Relationship. Does this tool leave more room for human connection, or quietly close it down?
- Development. Was it designed for the age and stage of the children using it, or adapted for them after the fact?
- Thinking. Does it extend children's thinking after understanding has formed, or remove the struggle through which understanding forms?
- Truth. Does its use build the habit of checking, or train children to accept fluent answers?
- Honesty. Are children told plainly that it is a machine, what it does with their data, and what is uncertain about it?

Two problems, not one

Alongside these, schools need to separate two problems that most policies treat as one. Academic misuse concerns how students produce assessed work and belongs with assessment design. Companion attachment concerns how students meet emotional and social needs and belongs with pastoral care and safeguarding (Penfold, 2026) [153]. Training pastoral

staff on the second is as important as updating integrity policy for the first.

What to do

For school leaders. Bring AI into pastoral and safeguarding training, not only integrity policy. Treat a student's use of AI for emotional support as a reason for a caring conversation, not a disciplinary one. Apply the five safeguards before adopting tools for younger children. Protect time and spaces for unmediated attention and relationships.

For teachers. Decide openly whether and when to tell students that feedback was produced with AI, and notice how it's received. Build slow, sceptical use of AI into lessons, especially with younger pupils. Teach the habit of asking how we know something is true.

For policymakers. Regulate companion chatbots for minors in ways that combine disclosure, crisis protocols and age-appropriate design with a role for parents. Fund research that follows children's AI use and well-being over time, since current evidence is mostly cross-sectional.

What this part can't tell you

Much of the evidence is cross-sectional or based on self-report and can't establish cause. The largest surveys come from the United States and Canada. The experimental evidence on cognitive offloading comes mainly from older students and adults, and its application to younger children is reasoned rather than tested. Several figures in the original newsletters were corrected or omitted because they couldn't be traced to primary sources. Regulatory proposals were still changing at the time of writing.

A closing thought

Children are growing up with machines that talk back. Most will use them sensibly most of the time, and many will benefit. But AI also reaches into the parts of childhood that schools care about most: relationships, attention, judgement and the sense of being known. Schools that treat AI only as an integrity problem will miss the student at lunch talking to a chatbot

instead of a classmate. The task isn't to keep children away from AI. It's to make sure it leaves room for the human relationships and the effortful thinking through which children grow, and to be honest with them about what they are talking to.

THINK IT THROUGH

Reflect.

- Where in a child's day does AI quietly replace a human conversation?
- What would you tell a child plainly about the machine they are talking to?
- If feedback on a child's work was written with AI, would you tell them? Why or why not?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- Relationship: it leaves more room for human connection, not less.
- Development: it was designed for the age and stage of the children, not adapted afterwards.
- Thinking: it extends thinking after understanding has formed, rather than removing the struggle.
- Truth: it builds the habit of checking and does not train children to accept fluent answers.
- Honesty: children are told plainly that it is a machine, what it does with their data and what is uncertain.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

PART FOUR

What Counts as AI Literacy

> Based on the paper "What Counts as AI Literacy? A Comparative Analysis of International Frameworks and System-Level Policies, 2024–2026" by Paul Penfold (Proceedings of Future Science Archives, 2026).

Part Three asked what children are owed. This part asks what young people should be able to do. It compares the main international definitions of AI literacy, looks at three school systems that acted on them, and offers a model that schools can use.

CHAPTER 18

Everyone agrees it matters, nobody agrees what it is

Ask a room of teachers whether AI literacy matters and almost every hand goes up. Ask what it covers, or how you would know whether a student had it, and the agreement falls apart. Schools have been trying to answer that second question, often from scratch, while their students' use of AI ran ahead of them.

A 2026 global survey of more than 45,000 people in higher education found that 88% of students and 77% of faculty were already using AI, but only 29% of students thought their instructors were equipped to guide them (Digital Education Council, 2026) [36]. Stanford's 2026 AI Index found that four in five K–16 students used generative AI for schoolwork, while only 6% of teachers knew whether their school had a policy on it (Stanford HAI, 2026) [175]. Chapter 32 returns to these figures when it looks at teachers.

Into that gap came a wave of frameworks. Academic work on AI literacy goes back at least to Long and Magerko (2020) [71], who proposed 17 competencies and a set of design considerations, and to Ng et al. (2021) [81], whose review grouped the field around knowing and understanding AI, using and applying it, evaluating and creating with it, and AI ethics. What changed after 2024 is that intergovernmental and international bodies began issuing definitions with real policy weight, and one of them, the OECD, began building an international test.

Three questions run through this part:

- Where do the main international AI literacy frameworks agree, and where do they differ?
- What do early system-level policies show about putting those definitions into practice?

- What should schools and curriculum designers take from both?

How the comparison was done

This is a comparative document analysis. I examined five frameworks published or finalised between 2024 and 2026, chosen because each was issued by an international body or widely used network and each sets out a structured definition of AI literacy for education. I also looked at three system-level policies that came to prominence in 2026: China's national plans for AI education, Boston Public Schools' high school AI literacy programme, and New York City's moratorium on student-facing generative AI from kindergarten to grade 8.

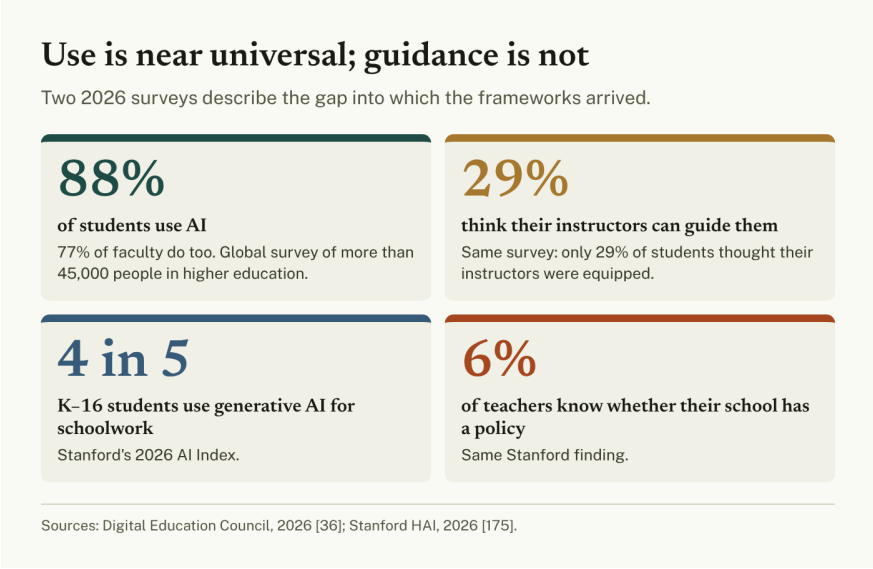


Figure: The usage and guidance gap (Digital Education Council, 2026; Stanford HAI, 2026).

The analysis began as a series of newsletters I wrote for school leaders (Penfold, 2026) [150, 157, 130, 109, 126, 106, 103, 120]. Each framework and policy was first read and discussed there as it was published. For this part, I went back to the primary documents and news sources and

checked the claims against them. I coded each framework for its definition, its structure, whether it describes progression, how it treats ethics and human agency, whether it addresses creation and design, and whether it says anything about assessment.

CHAPTER 19

Five frameworks side by side

UNESCO: a human-centred competency framework

UNESCO's student framework sets out 12 competencies across four dimensions: a human-centred mindset, the ethics of AI, AI techniques and applications, and AI system design. Each runs across three levels of progression: understand, apply and create (UNESCO, 2024) [185]. The companion framework for teachers adds dimensions on AI pedagogy and on AI for professional development (UNESCO, 2024) [186]. UNESCO is explicit that AI should serve people and remain under human control, and its framework treats ethics as a strand in its own right rather than a footnote.

Digital Education Council: general and domain literacy

The Digital Education Council (DEC) framework defines AI literacy as the knowledge and skills needed to understand, interact with and critically assess AI. It has five dimensions: understanding AI and data; critical thinking and judgement; ethical and responsible use; human-centricity, emotional intelligence and creativity; and domain expertise (Digital Education Council, 2025) [34]. The first four are general and apply to everyone. The fifth is specific to a person's subject or profession. Each dimension runs across three levels. Critical thinking and judgement, for example, moves from questioning AI output, to evaluating it, to challenging it.

Two features stand out. The split between general and domain literacy is something most other frameworks leave out. And DEC sets different target profiles for staff and students: faculty are expected to reach level 2 across the general dimensions, students a level 1 baseline. The framework was built with universities in mind, and its domain dimension needs real translation work before it fits a primary classroom, where a teacher's "do-

main" is as much child development as any single subject (Penfold, 2026) [150].

European Commission and OECD: the AILit Framework

The AILit Framework, finalised in 2026 after a 2025 draft, organises AI literacy into four domains: Engage with AI, Create with AI, Manage AI and Shape AI (European Commission & OECD, 2026) [43]. Each competence combines knowledge, skills and attitudes, and the domains are presented in sequence to suggest a possible learning path. It is aimed at primary and secondary education, is designed to work across subjects, and contributes directly to the PISA 2029 assessment. Its inclusion of Shape AI, which covers understanding how AI systems are designed and how society can influence them, makes it the framework most willing to treat young people as future builders and governors of AI, not just users.

OECD PISA 2029: literacy as something measured

The PISA 2029 Media and Artificial Intelligence Literacy (MAIL) framework is the first attempt to measure AI and media literacy among 15-year-olds at international scale. Its first draft, dated January 2026, defines MAIL as the competences needed to engage effectively, ethically and responsibly with digital content, media platforms and AI systems (OECD, 2026) [91]. It names five competences: access and use; analyse and evaluate; participate and collaborate; create; and reflect and act ethically and responsibly. The last of these sits beneath the others rather than alongside them. Tasks are set in five contexts: relationships, learning, entertainment, persuasion and citizenship.

What makes MAIL distinctive is its format. The example task places students in a simulated desktop, with a chat tool, browser and search engine, and asks them to verify a news story before publication, with sources of varying credibility. That's assessment of performance, not of recall. The OECD is explicit that the draft will change after piloting and a field trial (Penfold, 2026) [157].

Council of Europe: protecting social agency

The Council of Europe's recommendation to its 46 member states, adopted on 2 September 2026, sets out a three-dimensional model: a human dimension, a technological dimension and a practical dimension (Council of Europe, 2026) [27]. It singles out the human dimension, covering ethics, psychology, culture and impact on human rights, democracy and the rule of law, as deserving particular attention. It is blunt that AI systems "do not think, understand, reason or learn in the way humans do", and that people should not hand over their agency, judgement or responsibility to them.

It frames the purpose of AI literacy as protecting social agency, the capacity to act autonomously and meaningfully in society. It asks for AI literacy to be woven across civics, digital and media literacy and STEM rather than taught as a separate subject, and for honest communication with learners and families about when AI is being used in their education. Like all Committee of Ministers recommendations, it isn't legally binding (Penfold, 2026) [130].

At a glance

- UNESCO students (2024) [185]: 4 dimensions and 12 competencies; progression through understand, apply and create; distinctive for its human-centred mindset and AI system design.
- DEC (2025) [34]: 5 dimensions; 3 levels, with different targets for staff and students; distinctive for general plus domain-specific literacy.
- EC–OECD AILit (2026) [43]: 4 domains and 19 competences; domains presented as a sequence; distinctive for shaping AI and for combining knowledge, skills and attitudes.
- OECD PISA 2029 MAIL (draft, 2026) [91]: 5 competences and 5 contexts; a single age group (15); distinctive for performance-based, simulated tasks.
- Council of Europe (2026) [27]: 3 dimensions; progression not specified; distinctive for the human dimension, social agency and cross-curricular teaching.

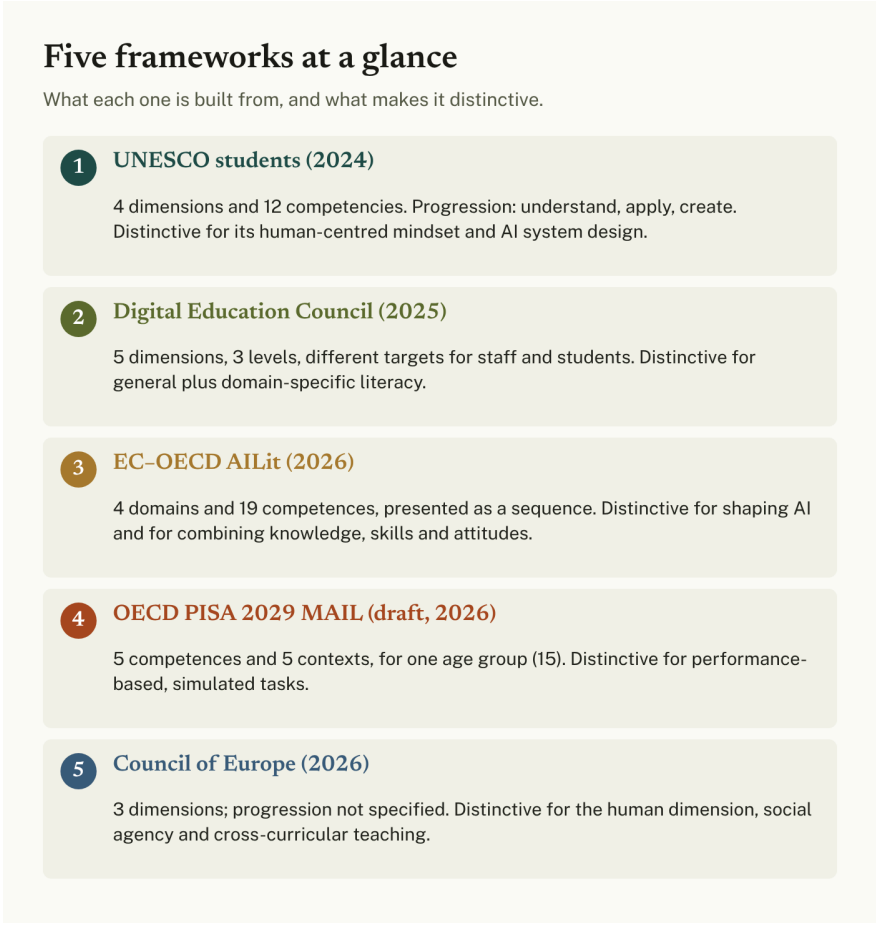


Figure: The five frameworks compared in Chapter 19.

CHAPTER 20

Four points of agreement, five real differences

Where the frameworks agree

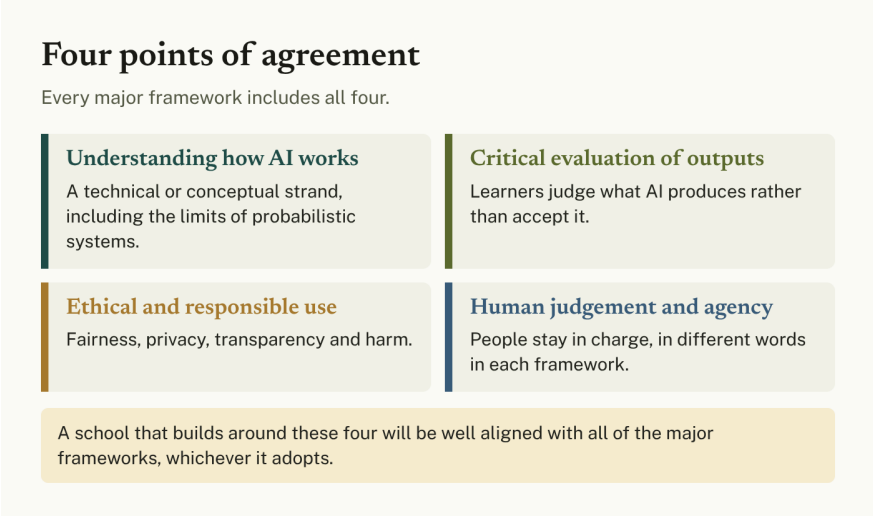


Figure: Where the frameworks agree.

Despite different origins, the frameworks converge on four elements.

- Understanding how AI works. Every framework includes a technical or conceptual strand: how data, models and outputs relate, and what the limits of probabilistic systems are. None of them treats this as optional.
- Critical evaluation of outputs. Every framework asks learners to judge what AI produces rather than accept it. DEC makes this a dimension of its own. MAIL builds half of its assessment time around analysing, evaluating and creating.

- Ethical and responsible use. Every framework includes ethics, covering fairness, privacy, transparency and harm. In MAIL and the Council of Europe model, ethics is a layer that runs beneath everything else, not a separate topic.
- Human judgement and agency. Every framework insists, in different words, that people stay in charge: UNESCO's human-centred mindset, DEC's human-centricity, the Council of Europe's social agency, AILit's attention to attitudes and to shaping AI.

The practical implication is reassuring. A school that builds its AI literacy work around these four elements will be well aligned with all of the major frameworks, whichever it eventually adopts.

Where they differ

The differences matter more for curriculum design than the agreements do.

How much weight goes on the human dimension. The Council of Europe gives the human dimension priority. Most school practice does the opposite, concentrating on tool use, prompting and spotting deepfakes, which the recommendation would place in the practical dimension (Penfold, 2026) [130]. A curriculum that starts and ends with tool training has covered the easiest third and skipped the two parts that matter more.

Whether learners create and shape AI. UNESCO's AI system design dimension and AILit's Shape AI domain treat learners as potential designers of AI and citizens who can influence it. DEC and the Council of Europe say much less about this. The difference reflects a real choice about whether AI literacy is a form of consumer protection or a form of participation.

General or domain-specific literacy. DEC alone separates general literacy from domain expertise. This matters because AI literacy looks different for a historian, a nurse and a primary teacher, even when the underlying judgement and ethics are the same. The other frameworks assume general literacy is enough, or leave subject application to teachers.



Figure: Five real differences between the frameworks.

How literacy develops with age. UNESCO and DEC describe levels, and AILit suggests a sequence, but none sets out a full developmental progression from early childhood to adulthood. National policy has moved faster here. China's plans tier AI education by age, with younger children experiencing AI in class and older students working up to creative projects with advanced tools (Bloomberg, 2026 [16]; Penfold, 2026 [109]). That's a claim about what AI literacy is at each stage, not just a rollout timetable.

Whether choosing not to use AI is a competence. The frameworks mostly describe competence as the ability to use AI well. An emerging strand of practice argues that the ability to decide not to use AI, for good reasons, should be named as a competence too. Teachers and school leaders at a 2026 UNESCO and European Union workshop made this case for

teachers specifically (Penfold, 2026) [126]. A teacher who keeps a chatbot out of a writing unit because it would undermine the thinking she's trying to build is showing exactly the judgement AI literacy is meant to produce. Systems that measure success only by adoption can't tell her apart from someone who has been left behind.

CHAPTER 21

China, Boston and New York City

Three policies from 2026 show how differently school systems have acted on the same broad idea.

China: national mandate, tiered by age

China's Ministry of Education set out its approach in a 2024 notice aiming to make AI education "basically universal" in primary and secondary schools by 2030, tiered so that younger pupils experience AI while older students work towards creative projects with advanced tools (Ministry of Education of the People's Republic of China, 2024) [73]. Beijing introduced compulsory AI lessons in primary and middle schools from autumn 2025 (NPR, 2026) [83], and in June 2026 the State Council issued a five-year plan to promote AI education throughout all stages of education (Bloomberg, 2026) [16].

The design is tiered by age and sits inside long-term planning. Its strength is coherence and scale. Its risk is speed: a mandate that runs ahead of settled pedagogy and trained teachers can produce box-ticking lessons rather than real judgement (Penfold, 2026) [109].

Boston: literacy for every high school graduate

Boston Public Schools launched an AI literacy programme in its high schools in September 2026, in partnership with the University of Massachusetts Boston, with teachers from each high school trained as AI ambassadors and district benchmarks for grades 7 to 12 under development (WBUR, 2026) [198]. It was widely reported as making AI fluency a graduation requirement, though the district's own framing has been closer to a proficiency goal for every graduate. The curriculum covers how AI systems work, their limits and biases, how to verify their outputs, and how

to use AI as a thinking tool rather than an answer machine (Penfold, 2026) [106].

New York City: a pause for younger learners

In September 2026, New York City announced a one-year moratorium on student-facing generative AI from kindergarten to grade 8, roughly 600,000 pupils, after two years of encouraging classroom use (CNN, 2026 [24]; StateScoop, 2026 [176]). Companion chatbots were barred at every grade. Teachers could still use AI for planning and administration but not for grading. A limited number of supervised high school pilots continued, alongside new mandatory AI critical-thinking courses, and a named coalition was given a year to review the evidence (Penfold, 2026) [103].

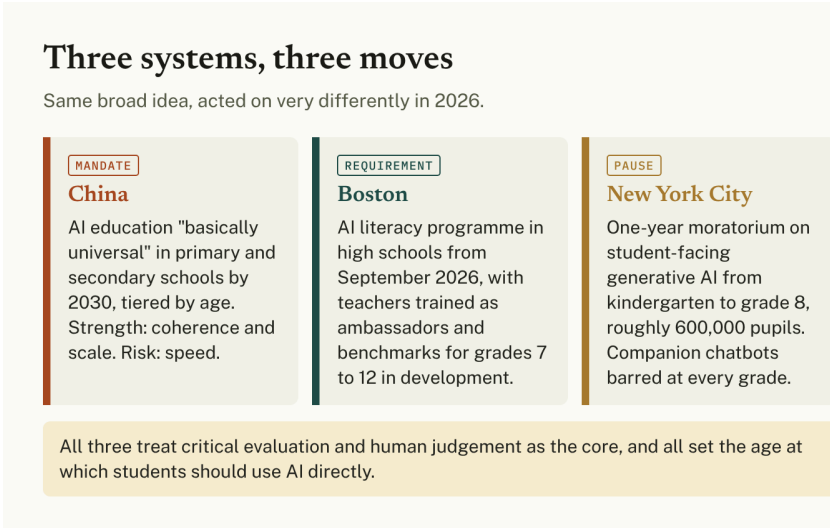


Figure: China, Boston and New York City (Chapter 21).

What the three share

On the surface, these look like opposites: one system mandating, one requiring, one pausing. Underneath, they share more than it seems. All three treat critical evaluation and human judgement as the core of AI literacy. All three, in different ways, set the age at which students should use AI

directly. And all three face the same two problems, which no framework solves on its own.

CHAPTER 22

What the definitions leave out, and a model for schools

Teacher capacity

Every policy depends on teachers who are themselves AI literate, and the data suggests most are not yet. The survey evidence in Chapter 18 points to the same gap: students' use is near universal, but confidence in teachers' ability to guide it is low, and policy is often invisible to the people meant to carry it out (Digital Education Council, 2026 [36]; Stanford HAI, 2026 [175]). China's attention to teacher standards, Boston's ambassador model and DEC's higher target level for staff than students are all attempts to address this. Teacher capacity is better treated as the first stage of any AI literacy programme, not a parallel workstream. You can't teach a child to weigh up something you're avoiding yourself (Penfold, 2026) [120]. Part Seven takes this up.

Assessment that measures doing, not knowing

The hardest question for any AI literacy requirement is how to tell a student who can verify a convincing AI answer from one who can only describe verification. A multiple-choice question on hallucinations tests whether a student can define one, not what they'll do when a fluent, confident, wrong answer appears. This question was put directly in response to Boston's announcement, because a credential is only useful if employers can trust it (Penfold, 2026) [106].

The strongest existing basis for an answer comes from outside AI. Wineburg and McGrew's (2019) [200] research on how professional fact-checkers evaluate online information found that experts leave a page almost immediately to check who is behind it and what other sources say, a practice they call lateral reading. Unlike general advice to "think critically",

this is a specific, teachable and testable method. PISA 2029's simulated verification task follows the same logic.

Performance-based tasks of this kind can be built for any subject. Students are given a fluent AI-generated text containing a fabricated citation, a distorted statistic, an unsupported causal claim and one accurate claim, and must decide what to trust. I've developed one such teaching prototype (Penfold, 2026) [106]. It isn't a validated instrument, but it shows what the shift from describing a skill to practising it looks like.

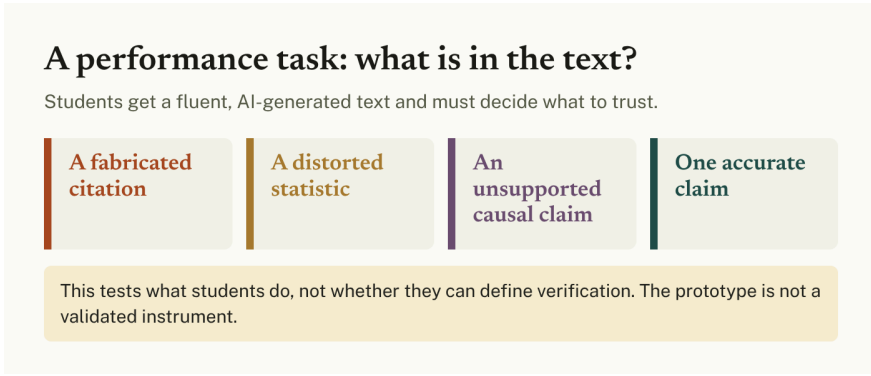


Figure: A performance-based verification task (Chapter 22).

A synthesis model

Drawing the frameworks and policies together, I propose that a school AI literacy programme needs five strands, a developmental progression, and assessment that matches each strand.

- Understand. How AI systems work, what data they depend on, and what they are not. (Roots: all five frameworks.)
- Evaluate. Question, verify and challenge AI outputs using methods such as lateral reading. (Roots: all five; DEC and MAIL especially.)
- Use well. Apply AI to real tasks in a subject or field, including knowing when it helps learning and when it doesn't. (Roots: DEC domain expertise; AILit Engage and Create.)

- Act ethically. Fairness, privacy, transparency, accountability, and honesty about use. (Roots: all five; the Council of Europe's human dimension.)
- Decide and shape. Keep human agency: choose when not to use AI, and understand how society can shape AI. (Roots: the Council of Europe's social agency; AILit Shape; UNESCO system design.)

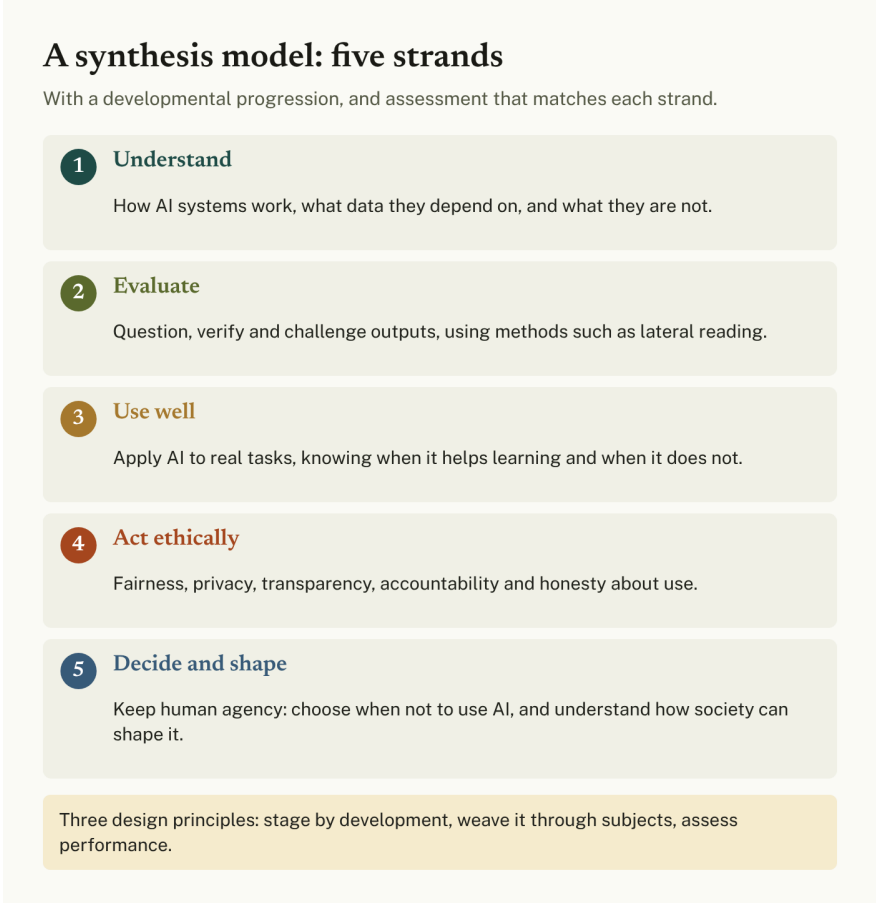


Figure: The five strands of the synthesis model (Chapter 22).

Three design principles follow.

- Stage by development, not by when the software arrived. Younger children should experience and question AI with an adult, while older students can build and create with it. China's tiered model and New York City's age-based pause both reflect this, from opposite directions.
- Weave it through subjects. The Council of Europe and AILit both favour cross-curricular integration. Evaluation and ethics belong in history, science and English as much as in computing.
- Assess performance. Use tasks that require students to verify, judge and decide, not just define.

What to do

For curriculum designers. Build around the four elements on which every framework agrees, then make explicit choices on the five areas where they differ. Hold the curriculum loosely: the PISA framework is still a draft, and systems that build too tightly around today's version may need to rebuild after 2029 (Penfold, 2026) [157].

For school leaders. Start with teachers. Audit staff and student literacy against a named framework rather than guessing from anecdote. Give teachers room for a well-reasoned "no", and measure judgement rather than adoption. Whatever you decide about classroom AI, name who reviews the evidence and when, and tell your community (Penfold, 2026) [103].

For assessment bodies and credentialing systems. If AI literacy becomes a requirement for graduation or certification, invest in performance-based assessment before scaling the requirement. A credential that measures knowledge about verification won't tell employers what they need to know.

What this part can't tell you

This analysis covers frameworks and policies published in English or reported in English-language sources, and it draws on a newsletter series that selected topics for their relevance to school leaders, not through a systematic search. The PISA framework is a draft and will change. The

status of Boston's requirement was still being settled at the time of writing, and reports of China's plans rely partly on secondary coverage. These policies are moving quickly, and some of what is described here may already have been revised. Finally, I work in school leadership and curriculum design, which shapes the questions I ask.

A closing thought

The international effort to define AI literacy has been more successful than it might appear. The major frameworks agree on what matters most: understanding the technology, evaluating what it produces, using it ethically and keeping human judgement in charge. Their differences, on the human dimension, on creation, on domain knowledge, on development and on the right to say no, are real choices that schools should make deliberately rather than by default. The harder problems now lie outside the definitions, in teachers who are confident enough to teach AI literacy and in assessment that can tell knowing from doing.

THINK IT THROUGH

Reflect.

- Which of the five strands is weakest in your setting, and why?
- Where do you think teachers should have room for a well-reasoned 'no'?
- How would you tell the difference between a student who knows about verification and one who can do it?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- Understand: students learn how AI systems work, what data they depend on and what they are not.
- Evaluate: students practise questioning and verifying AI output, for example by lateral reading.
- Use well: students apply AI to real tasks in a subject, including knowing when it does not help learning.
- Act ethically: fairness, privacy, transparency and honesty about use are taught and discussed.
- Decide and shape: students are free to choose not to use AI and understand how society can shape it.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

PART FIVE

When the Artefact Is No Longer Evidence

> Based on the paper "When the Artefact Is No Longer Evidence: Assessment Validity and Academic Integrity in the Age of Generative AI" by Paul Penfold (Proceedings of Future Science Archives, 2026).

Part Four asked what young people should be able to do with AI. This part asks how we would ever know. It looks at why the essay and the take-home task have stopped working as proof, why detection can't repair them, and what schools can build in their place.

CHAPTER 23

The quiet assumption behind most assessment

A lot of school assessment rests on an assumption nobody says out loud: that producing the work proves you can do the thinking. Write the essay and you must have understood the argument. Solve the problems at home and you must have grasped the method. For a long time the assumption held, because producing the work required the thinking. It no longer does.

This is often described as a cheating crisis, and the first reflex of many institutions has been to reach for detectors, bans and locked-down conditions. I think the reflex misreads the problem. The essay was always a proxy for the understanding behind it. If AI can now produce the proxy without the understanding, the proxy has stopped working, and no amount of policing will bring it back. The more useful question, put plainly by Dawson and colleagues (2024) [31], is about validity rather than cheating: does this assessment still let us make a trustworthy judgement about what a learner knows and can do?

Three questions guide this part:

- What does the evidence say about how students use AI in assessed work, and how teachers see it?
- Why is detection an inadequate response?
- What assessment designs give more trustworthy evidence, and what are their risks?

Where the material comes from

This is a practitioner-informed synthesis. Its starting point is a set of newsletter articles on assessment and integrity that I wrote for school leaders between June and October 2026 (Penfold, 2026) [113, 128, 159,

105, 96, 114]. They include accounts of assessment redesign at an international school in Chiang Mai, Thailand, where I worked, and summaries of reports from UNESCO, the OECD and the World Economic Forum. For this book I went back to the primary sources and checked them, and corrected figures where the newsletters had misreported them. The school practice is described as I reported it. It hasn't been independently evaluated.

CHAPTER 24

How big is the problem?

How students use AI

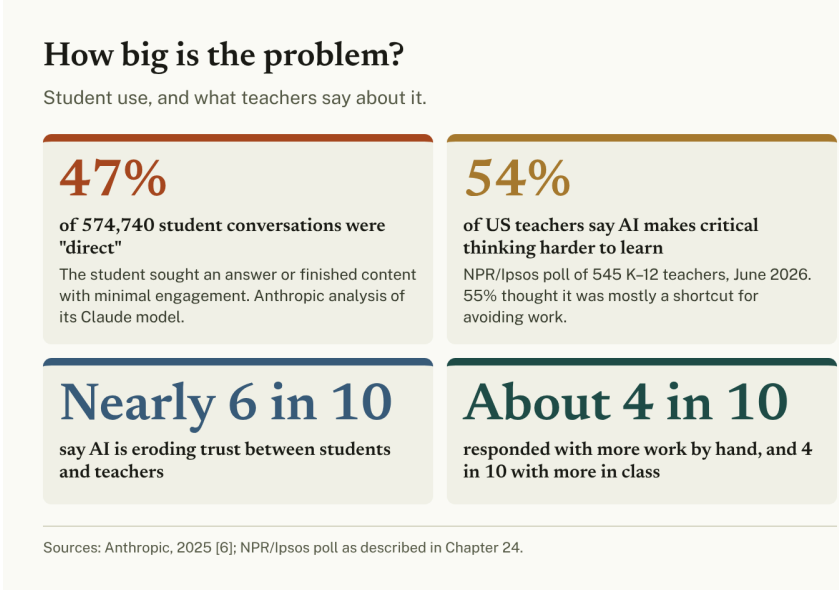


Figure: Student use of AI and teachers' views (Chapter 24).

Anthropic analysed 574,740 anonymised student conversations with its Claude model. Nearly half, about 47%, were "direct" conversations, in which the student sought an answer or finished content with minimal engagement. Some students explicitly asked for help avoiding plagiarism detection (Anthropic, 2025) [6]. Not all direct use is dishonest. But the pattern shows how easily assessed work can now be produced without the thinking it is meant to evidence.

How teachers see it

A nationally representative NPR/Ipsos poll of 545 US K–12 teachers, published in June 2026, found that 54% believed AI makes it harder for students to learn critical thinking skills, and 55% thought it was mostly a shortcut for avoiding work. Nearly six in ten said AI was eroding trust between students and teachers. About four in ten had responded by requiring more work to be done by hand, and four in ten by requiring more work to be done in class. Among teachers whose school provided AI software, only 35% said there was a formal policy on teacher use (NPR, 2026) [84]. Teachers are already redesigning assessment, mostly on their own and without guidance.

From shortcuts to automation

In February 2026, a tool called Einstein was released that could log into the Canvas learning platform, watch lectures, write essays and submit coursework from start to finish (Inside Higher Ed, 2026) [56]. Cease-and-desist letters took it offline within weeks, but Canvas's owner then launched its own agentic AI feature, framed as support for teaching (Inside Higher Ed, 2026) [57].

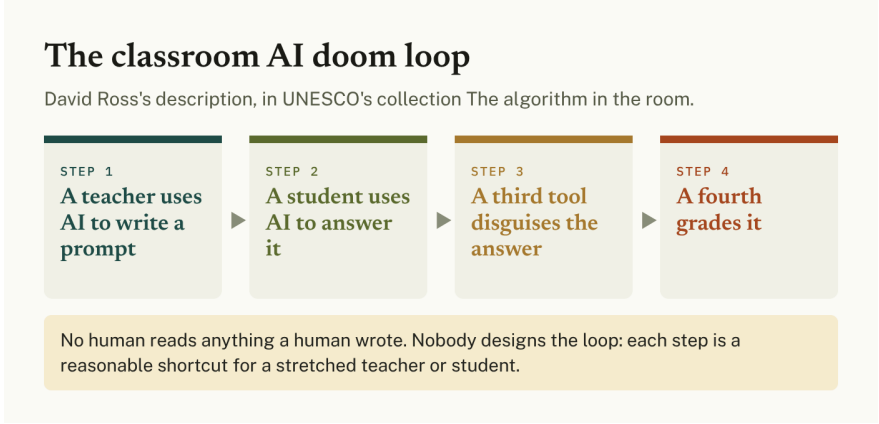


Figure: The classroom AI doom loop (Chapter 24).

David Ross, writing in UNESCO's collection The algorithm in the room, describes where this leads: the "classroom AI doom loop". A teacher uses AI to write a prompt, a student uses AI to answer it, a third tool disguises the answer, a fourth grades it, and no human reads anything a human wrote (UNESCO, 2026 [189]; Penfold, 2026 [128]). Nobody designs the loop. Each step is a reasonable shortcut for a stretched teacher or student, which is why it's so hard to stop.

CHAPTER 25

Why detection is the wrong answer

Detectors are unreliable

Weber-Wulff and colleagues (2023) [199] tested fourteen AI detection tools, including commercial products, and concluded that they were neither accurate nor reliable. Accuracy fell further when AI-generated text had been lightly edited or paraphrased, which is exactly what a student trying to avoid detection would do. A detector whose result can be changed by a paraphrasing tool can't carry the weight of a misconduct finding.

Detectors are biased

Liang and colleagues (2023) [68] found that GPT detectors consistently misclassified writing by non-native English speakers as AI-generated, while classifying writing by native speakers correctly. The likely reason is that detectors treat simpler, more predictable language as a sign of machine authorship, and second-language writers often write that way. For international and multilingual schools, this turns a technical weakness into an equity and safeguarding problem. An integrity policy that rests on detection will fall hardest on the students with the least power to contest it (Penfold, 2026) [159]. Part Ten comes back to what this means for fairness.

Surveillance damages the relationship it's meant to protect

The World Economic Forum's 2026 report on education readiness for AI warns that monitoring tools risk turning the student–teacher relationship from one built on trust and development into one premised on monitoring and suspicion (World Economic Forum, 2026 [205]; Penfold, 2026

[105]). Answering a trust problem with surveillance moves the trust problem somewhere else. It doesn't solve it.



Figure: Three problems with AI detection (Chapter 25).

CHAPTER 26

From integrity to validity

Grades as a social contract

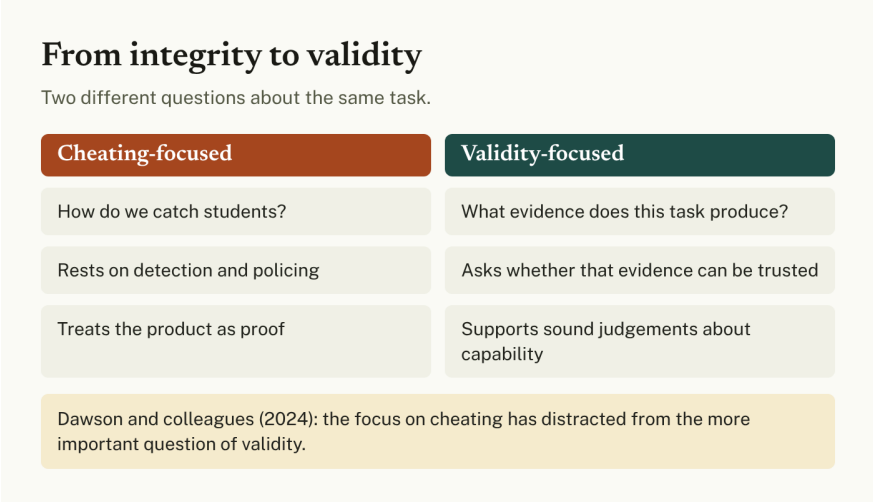


Figure: Cheating versus validity (Chapter 26).

The same WEF report reframes academic integrity as being clear about what students did themselves, what AI helped with, and whether the assignment still lines up with its intended learning outcome. It describes grades as a kind of social contract: institutions attest that a grade reflects real understanding, and employers, universities and the public act on that attestation. When the contract breaks, the damage spreads beyond the classroom to employer confidence and belief in merit (World Economic Forum, 2026) [205].

A different question

Dawson and colleagues (2024) [31] argue that the focus on cheating has distracted from the more important question of validity: whether assessment supports sound judgements about student capability. A cheating-focused response asks how to catch students. A validity-focused response asks what evidence a task produces and whether that evidence can be trusted.

Australia's higher education regulator took a similar line in guidance commissioned in 2023, arguing that trustworthy judgements about learning in a time of AI need multiple, inclusive and contextualised approaches to assessment rather than a single secure point of measurement (Lodge et al., 2023) [70].

What the learning research adds

The learning research makes the same point from another side. Chapter 10 describes a field experiment with nearly a thousand high school mathematics students, in which those who practised with an unrestricted chatbot performed far better during practice but worse than a no-AI control group on an exam taken without it (Bastani et al., 2025) [11]. Their practice scores were real, but they weren't valid evidence of what the students could do alone. Any assessment that measures performance with AI available, without also checking what the learner can do without it, risks recording the tool's capability as the student's.

The single sitting is fragile too

The 2026 examination season showed that the problem isn't only coursework. In Mexico, around 58,000 applicants to the national university had to retake an entrance exam, held online for the first time, after abnormally high marks raised suspicions of widespread cheating (CNN, 2026 [25]; Solomon & Holmes, 2026 [171]). In Portugal, errors in a newly digitised exam marking system set off a political crisis. In India, a paper leak in a national medical entrance examination led to mass protests and the resignation of the education minister. Thomas Lancaster of Imperial College

London described how cheating has moved from hidden notes to earpieces, hidden cameras and AI tools that supply answers in real time (Solomon & Holmes, 2026) [171].

Denmark responded by requiring upper secondary students to give an oral defence of their written essays from the new school year. The academic integrity researcher Sarah Eaton argued that, because AI can generate a perfect essay in seconds, the essay is no longer a viable way to demonstrate learning and should be allowed to "die with dignity" (Solomon & Holmes, 2026) [171]. Her wider point is that when one examination decides a young person's future, the pressure on students and the incentive to cheat both become extreme.

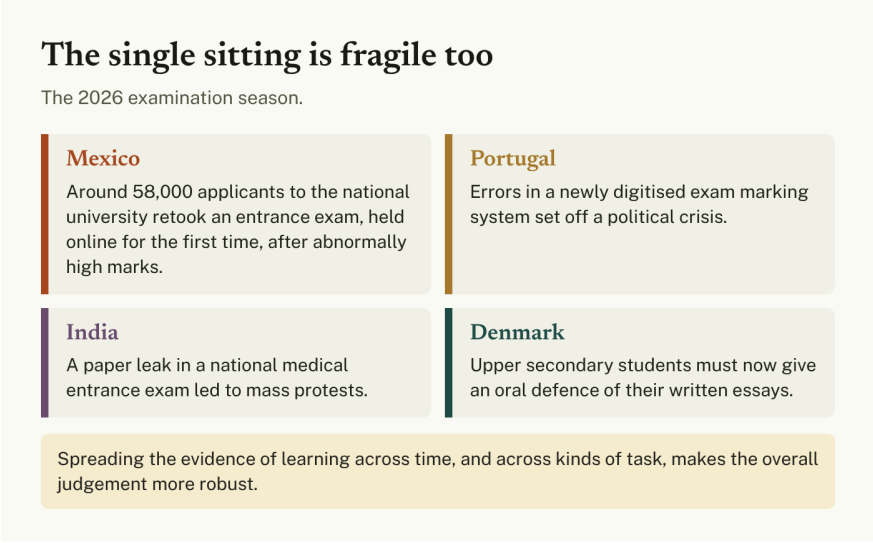


Figure: Examination problems in 2026 and Denmark's response (Chapter 26).

The lesson for school systems is that a single sitting at the end of a course is now a weak point in its own right, exposed to leaks, marking errors and real-time AI. Spreading the evidence of learning across time, and across different kinds of task, makes the overall judgement more robust (Penfold, 2026) [96].

CHAPTER 27

What schools can build, and the risks of building it

Six moves from one school



Figure: Six assessment moves from one school (Chapter 27).



Figure: Four prompts for a conversation (Chapter 27).

At the international school in Chiang Mai described in my newsletters, secondary teachers adopted six assessment moves across subjects (Penfold, 2026) [113]. None makes cheating impossible. The aim is to make honest work the easier path, and to require anyone faking it to keep up the deception across several kinds of evidence.

- Process portfolios. Learners submit planning notes, an annotated draft with teacher feedback, and a short reflection with the final piece, and the package is graded as a whole. In Year 10 English, the finished essay arrives with a visible history.
- Supervised components. Research happens at home, but the assessed section is written in class where a teacher can see it happen. In Year 9 History, sources were gathered across a week and the analysis written in a protected lesson.
- Personal and local reflection. Questions that only make sense if the learner was there, such as explaining how a soil test the class ran supported or contradicted a hypothesis. No AI attended that lesson.

- Oral defence. A short recorded viva in which learners explain their choices and answer follow-up questions, built into normal assessment so it doesn't feel like an extra hurdle. Year 11 science investigations included a five-minute viva on method and results.
- Iterative feedback. A draft is marked and the learner must fix specific, named weaknesses in the next version.
- Named-source constraints. Tasks are limited to sources the teacher provides, framed in a way a general model hasn't seen: "Using only Sources A, B and C from our lesson pack, evaluate..."

Conversations, not accusations

When a teacher suspects that AI wrote a piece of work, the school's approach was a conversation rather than an accusation. The conversation finds out what the learner understands, and is useful as feedback either way (Penfold, 2026) [113]. Four prompts do most of the work: "Talk me through how you structured this"; "What was the hardest part of this?"; "Put this sentence into your own words"; and, if needed, "Write the next paragraph now, on a related point, in the same style." The teacher keeps a brief note. One conversation is formative. A documented pattern is evidence.

Learners as producers

A further approach turns students from consumers of assessment into producers of learning materials. In a Year 10 co-creation model, students wrote concept chapters from the examination specification, built their own banks of exam-style questions and produced podcasts and infographics, all checked by their teacher against the syllabus (Penfold, 2026) [96]. This draws on the protégé effect: learners who prepare to teach material tend to work harder to understand it than learners who prepare only to be tested (Chase et al., 2009) [22]. It also leaves a running record of understanding that is hard to fake.

Old practice, newly relevant

None of this is new. Vocational and professional education has long relied on work-based and practical assessment, in which competence is observed in a real setting rather than inferred from a written product. Early-years traditions such as Reggio Emilia have long treated documentation of a child's thinking as evidence of learning. What AI has changed is that these approaches now matter in every subject, not just those where performance was always the point (Penfold, 2026) [114].

Equity and cultural risks

Moving away from written products isn't risk-free, and schools shouldn't assume the alternatives are fairer by default.

Oral and in-person assessment can disadvantage some learners. Students who are anxious, who are working in a second language, or who have speech, language or processing needs may show less of what they know in a viva than in writing. Oral defence needs preparation time, familiar formats and reasonable adjustments, and should sit alongside other evidence rather than replace it.



Figure: Equity and cultural risks (Chapter 27).

Assessment that rewards polish may reward sameness. Doshi and Hauser (2024) [38] found that generative AI raised the quality of individual creative writing while reducing the diversity of what a group produced. Agarwal and colleagues (2025) [3] found that AI writing suggestions nudged writers towards Western styles and reduced cultural nuance. Rubrics that reward fluency and finish may end up rewarding the voice of the model rather than the voice of the learner, so criteria may need to give more weight to originality, local knowledge and the learner's own reasoning (Penfold, 2026) [114]. Chapter 48 returns to this in the context of language and culture.

Unequal AI literacy creates unequal integrity risk. The WEF report notes that AI can give every student the kind of tutoring once available only to wealthier families, but that the line between getting help and outsourcing work is unclear, and students with less AI literacy are more likely to cross it without realising (World Economic Forum, 2026) [205]. Clear, task-by-task guidance on acceptable use is an equity measure as well as an integrity one.

An evidence chain

Drawing these threads together, I propose that every significant assessment task should be able to answer five questions. Together they form an evidence chain from purpose to judgement.

- Purpose. What capability is this task meant to show? (Weighing historical evidence, say, not producing a tidy essay.)
- AI role. Is AI excluded, allowed for support, or part of the task, and is that clear to learners? (AI allowed for feedback on a draft, not for the final analysis.)
- Process evidence. What evidence of the learner's thinking will exist besides the final product? (Planning notes, an annotated draft, a reflection.)
- Human judgement. Where does a teacher who knows the learner see the work and make a judgement? (A short viva or conversation on key choices.)

- Fairness. Who might this design disadvantage, and how is that managed? (Preparation time and adjustments for second-language learners.)

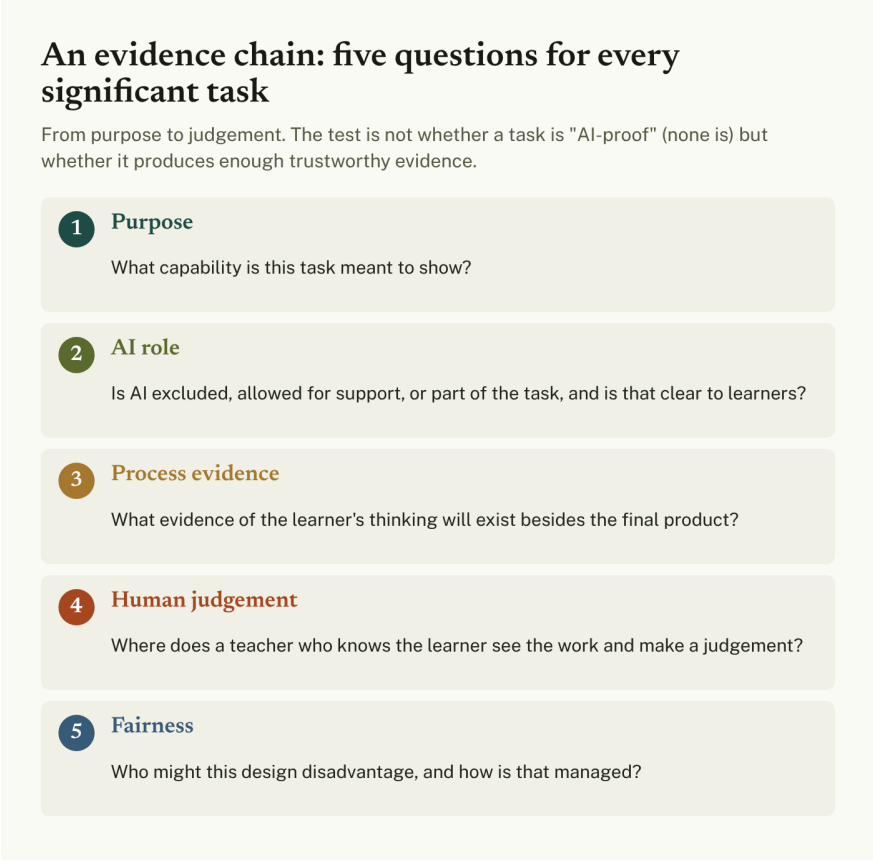


Figure: The five-question evidence chain (Chapter 27).

The framework doesn't ask whether a task is "AI-proof". No task is. It asks whether the task produces enough trustworthy evidence, from enough sources, for a teacher to make a sound judgement. It also gives schools a way to trace an assignment from the moment it's set to the moment it's graded, and mark every point where a human looks at the work. That's the

most direct test of whether a doom loop is already running (Penfold, 2026) [128].

What to do

For teachers. Decide what each task is for before deciding the AI rule. Sequence work so learners attempt first and bring AI in second. Keep a regular task done without tools, so you can see what learners can do alone. Treat suspected AI use as a conversation about understanding.

For school leaders. Don't build an integrity policy on detection tools, and review any existing use of them for bias against multilingual learners. Map the assessment pipeline and find where human judgement happens. Give teachers time and guidance to redesign tasks, and agree acceptable use with learners and families task by task.

For examination boards and systems. Reduce reliance on a single high-stakes sitting by combining it with evidence gathered over time. Build oral and practical components with adjustments for learners who need them. Invest in moderation of teacher judgement, since that judgement is becoming the main source of trustworthy evidence.

What this part can't tell you

The school practice here comes from my own account and hasn't been independently evaluated, so its effectiveness isn't established. The survey and news sources reflect particular countries and moments. Some figures in the original newsletters were corrected, and readers should rely on the cited primary sources. Research on oral and process-based assessment in the context of AI is still limited, especially in schools rather than universities.

A closing thought

Generative AI hasn't created a new problem for assessment so much as exposed an old one. Much school assessment measured production because production was easy to mark and usually required thinking. That link has gone. Detection tries to restore it by policing the product, and it fails, un-

fairly. The better path is to ask what each task is meant to show, collect evidence of the thinking as well as the product, and keep a teacher who knows the learner at the point of judgement. The machine writing the essay may be the push schools needed to assess what they always claimed to value.

THINK IT THROUGH

Reflect.

- What is this task meant to show about a learner, in one sentence?
- What would a teacher see in the learner's process, as well as the product?
- Who might your current method disadvantage, and what is the smallest adjustment that would help?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- Purpose: the task states what it is meant to show, and why.
- AI role: the brief says what AI use is allowed, task by task.
- Process evidence: drafts, notes, a log or a reflection show how the work came about.
- Human judgement: a teacher who knows the learner weighs the evidence.
- Fairness: the method works for every learner, including multilingual learners and those who find oral formats hard.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

PART SIX

The Broken Bottom Rung

> Based on the paper "The Broken Bottom Rung: Entry-Level Work, Durable Capability and What Schools Should Prepare Young People For" by Paul Penfold (Proceedings of Future Science Archives, 2026).

Part Five asked how we can trust evidence of learning. This part turns to what young people are learning for. It looks at what AI is doing to the first rung of working life, why that matters to schools, and how education might help rebuild it.

CHAPTER 28

The place where people learn to work

Most of what we hear about AI and work points one way: fewer jobs, smaller teams, doing more with less. The evidence is more varied than that, and more useful to educators. AI doesn't affect all work equally. Its effects are showing up first in particular tasks, particular occupations and, most importantly for schools, at a particular stage of working life.

Entry-level work has always been more than a job. It's how people learn to do the job at all. Junior analysts learn which numbers to trust. New developers learn how systems break. Trainee lawyers learn how rules, deadlines and human judgement collide. That learning happens through routine work: drafting, checking, summarising and tidying up. This is exactly the layer of work that AI now does well (Penfold, 2026) [142].

Three questions guide this part:

- What does recent evidence say about how AI is affecting work, especially entry-level work?
- What does that mean for the skills and pathways schools should prioritise?
- How should schools prepare young people for a labour market whose entry points are changing?

Where the material comes from

This part draws on newsletters on skills and the future of work that I wrote for school leaders in 2026 (Penfold, 2026) [142, 125, 144, 99, 154, 132, 108, 143, 149]. They summarised reports from the World Economic Forum, PwC, the International Labour Organization, the Digital Education Council and others. I checked the figures against primary sources and updated them where newer versions of a study had been released. The analysis uses the task-based approach to technological change

in labour economics, which treats jobs as bundles of tasks, some of which technology can perform and some of which it complements (Autor et al., 2003) [10].

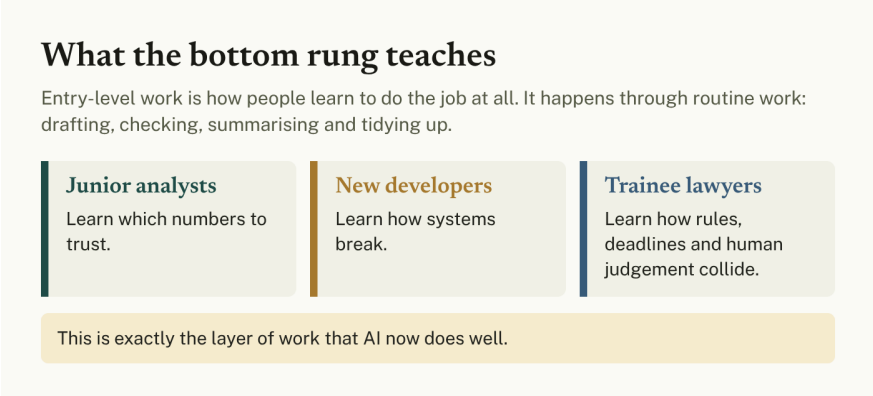


Figure: What entry-level work teaches (Chapter 28).

CHAPTER 29

What the labour market is signalling

Skills are shifting quickly

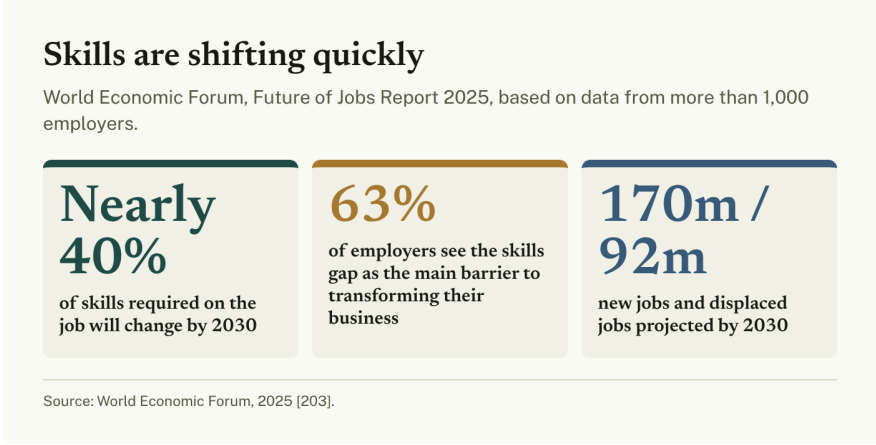


Figure: Projected skill change to 2030 (World Economic Forum, 2025).

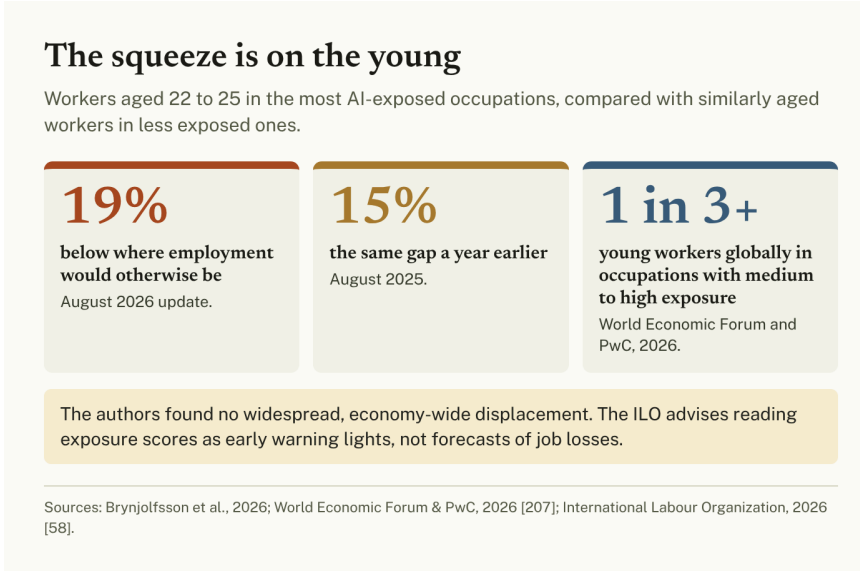


Figure: Early-career employment in AI-exposed occupations (Chapter 29).

The World Economic Forum's Future of Jobs Report 2025, based on data from more than 1,000 employers, projected that nearly 40% of the skills required on the job will change by 2030, and that 63% of employers already see the skills gap as the main barrier to transforming their businesses (World Economic Forum, 2025) [203]. Technology skills such as AI, big data, and networks and cybersecurity were expected to grow fastest. Human skills such as analytical thinking, creative thinking, resilience, flexibility and leadership stayed critical. The report projected 170 million new jobs and 92 million displaced by 2030, a net gain of 78 million, with significant growth in care and education roles as well as technology (Penfold, 2026 [125], 2026 [144]).

The squeeze is on the young

The most striking evidence concerns early-career workers. Using payroll data, Brynjolfsson, Chandar and Chen found that employment for workers aged 22 to 25 in the most AI-exposed occupations fell relative to their peers after generative AI became widespread, while older workers in the

same occupations didn't see the same decline. Their August 2026 update found employment for this group about 19% below where it would have been had it kept pace with similarly aged workers in less exposed occupations, up from 15% a year earlier. The authors found no widespread, economy-wide displacement (Brynjolfsson et al., 2026) [19]. Their estimates have changed between versions of the study, so treat the exact percentages with care and check the latest update.



Figure: Democratised and professionalised roles (PwC, 2026).

A 2026 report by the World Economic Forum and PwC found that more than one in three young workers globally are in occupations with medium to high exposure to AI-driven change in their tasks (World Economic Forum & PwC, 2026 [207]; Penfold, 2026 [142]). Entry-level workers surveyed by PwC said AI had made them more productive, but many also said they were working longer because of it.

Exposure is not the same as job loss

The International Labour Organization warns against reading AI "exposure" scores as forecasts of job losses. Exposure measures what AI could do to a set of tasks, not whether firms will automate them or how work will be reorganised. The ILO advises treating such measures as early warning lights, to be read alongside real data on hiring, wages and job moves (International Labour Organization, 2026 [58]; Penfold, 2026 [142]). It

also notes that newer, AI-specific measures reverse an older assumption: it's now higher-skilled cognitive work, in business, finance, computing and education, that scores as most exposed.

Growth, and a premium for skill

Exposure hasn't meant shrinking. PwC's 2026 Global AI Jobs Barometer, based on job advertisements across 27 countries, found that headcount at the companies most exposed to AI grew by 52% from 2018 levels, compared with 36% at the least exposed (PwC, 2026 [161]; Penfold, 2026 [99]). Workers with AI skills commanded an average wage premium of 62%, up from 57% the previous year, ranging from 16% in government to 118% in consumer markets.

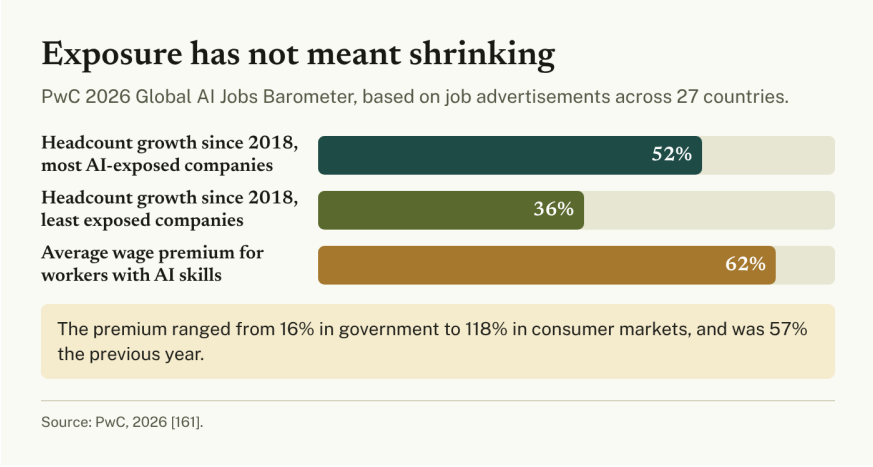


Figure: Growth and the wage premium for AI skills (PwC, 2026).

PwC also identified two paths in the jobs AI reshapes. In "democratised" roles, AI takes on the more skilled task and leaves the simpler one to a person, and such roles tend to shrink and pay less. In "professionalised" roles, AI takes the routine tasks and leaves the harder ones, such as interviewing, negotiating and judging, and such roles tend to grow and pay more (PwC, 2026) [161]. The jobs AI had reshaped also placed more

weight on distinctly human capabilities such as empathy, judgement and creativity (Penfold, 2026) [99].

CHAPTER 30

The expertise paradox

The Digital Education Council analysed tasks across eleven job families. It found that automation is densest in structured, rule-based work, while banking and healthcare keep the largest share of human-led tasks because accountability creates a ceiling: when someone must answer for a decision, a person stays in the loop (Digital Education Council, 2026 [37]; Penfold, 2026 [154]).

Its most important finding for education is what I call the expertise paradox: AI increases the value of expertise while making it harder to build.

The people best placed to judge where and how AI should be used are those with deep domain knowledge, and that knowledge has traditionally come from doing the entry-level work AI is now absorbing. If firms automate the routine layer and stop training juniors, they may look more efficient this year and find in a decade that no one understands how the work really behaves (Penfold, 2026) [142].

The Council also found that the management skills of coordinating, overseeing and evaluating work, which people used to grow into over years, are becoming early-career skills. A junior employee working alongside AI has to manage its output from the first day (Penfold, 2026) [154].

Some employers are choosing to protect the pipeline. The WEF and PwC framework argues that entry-level hiring should be a deliberate part of workforce planning, not a cost to cut when AI makes it look easy (World Economic Forum & PwC, 2026) [207]. For education, the implication is direct: if workplaces provide fewer routine learning opportunities, schools and universities will need to provide more authentic ones.



Figure: The expertise paradox (Chapter 30).

CHAPTER 31

Credentials, caring work and a three-layer model

The weakening signal

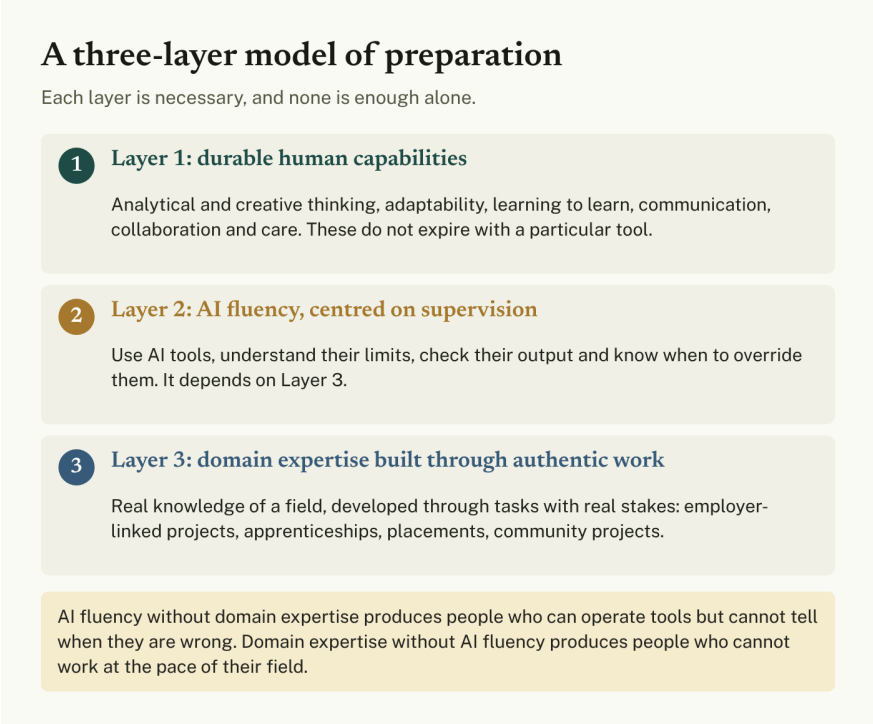


Figure: The three-layer model of preparation (Chapter 31).

Spence's (1973) [173] account of job market signalling explains why credentials carry value: they let employers infer qualities they can't observe directly. That only works while employers and graduates broadly agree on what a credential represents. The WEF's 2026 education readiness report

notes growing scepticism that academic credentials lead to economic opportunity, as skills change faster than qualification frameworks can track (World Economic Forum, 2026 [205]; Penfold, 2026 [132]). Goldin and Katz's (2008) [50] history of the "race between education and technology" shows that inequality widens whenever education falls behind technological change.

Rapid restructuring

Some systems are responding at scale. Between 2021 and 2025, Chinese universities revoked or suspended 12,200 undergraduate degree programmes and introduced 10,200 new ones. The cuts were concentrated in arts, humanities, foreign languages and management, and the new programmes were aligned with national development goals, including AI (South China Morning Post, 2026 [172]; Penfold, 2026 [108]). The scale is unmatched elsewhere, and so are the risks: fields whose value doesn't show up in employment statistics may be lost. But the underlying discipline of auditing what is taught against what the world now needs is one that other systems could borrow at their own pace.



Figure: Degree programmes cut and added in China, 2021 to 2025.

Caring and vocational work

The WEF projects significant growth in care roles such as nursing, as well as in education roles (World Economic Forum, 2025) [203]. As machines take over routine tasks, the work that remains leans harder on what machines can't do, and caring for another person is near the top of that list. Yet schools often steer their highest-achieving students towards academic routes and treat caring and vocational pathways as a fallback (Penfold, 2026) [143]. That's a moral mistake and, increasingly, an economic one.

Writers in UNESCO's collection make a parallel point about vocational education: craftsmanship rests on judgement developed through the hands over years, which AI can support but not shortcut (Haarsbjerg, 2026 [55]; Penfold, 2026 [149]).

A three-layer model of preparation

Drawing these strands together, I propose that schools prepare young people in three layers. Each is necessary, and none is enough alone.

- Layer 1: durable human capabilities. Analytical and creative thinking, adaptability, learning to learn, communication, collaboration and care. These don't expire with a particular tool, and employers consistently rank them highly (World Economic Forum, 2025) [203]. Schools often claim to value them while their timetables and assessments still reward content recall (Penfold, 2026) [125].
- Layer 2: AI fluency, centred on supervision. Not "learn to code" or "learn to prompt", both of which date quickly, but the ability to use AI tools, understand their limits, check their output and know when to override them. This is the early-career management skill that the Digital Education Council describes, and it depends on Layer 3 (Penfold, 2026 [142], 2026 [154]).
- Layer 3: domain expertise built through authentic work. Real knowledge of a field, developed through tasks with real stakes: employer-linked projects, apprenticeships, work placements, community projects and co-operative education. If the workplace's bottom rung is weaken-

ing, education needs to rebuild some of it, because judgement is built in real settings (Penfold, 2026 [142], 2026 [144]).

The layers depend on each other. AI fluency without domain expertise produces people who can operate tools but can't tell when they're wrong. Domain expertise without AI fluency produces people who can't work at the pace of their field. Both depend on durable capabilities that let a person keep learning as the tools and the market change.

What to do

For school leaders. Audit where the curriculum still rewards recall over reasoning, and give teachers time to teach durable capabilities explicitly. Build live links with employers through projects, mentoring and placements, especially for students without family networks, who depend most on school to build the bridge to work (Penfold, 2026) [144]. Present caring and vocational pathways as ambitious choices, not consolation prizes.

For curriculum designers and universities. Teach AI fluency as supervision and verification within subjects, not as a standalone tool course. Review programmes regularly against what graduates now need, with care not to discard fields whose value isn't captured by employment data.

For employers and policymakers. Treat entry-level hiring and training as an investment in the future workforce, not a cost to automate away. Recognise evidence of capability, such as portfolios and observed work, alongside traditional credentials.

What this part can't tell you

Much of the evidence comes from the United States and from advanced economies. Exposure measures describe potential task change, not outcomes. The strongest evidence on young workers comes from a single, repeatedly updated study using payroll data, and its estimates have changed between versions. Several figures in the original newsletters were updated or omitted where primary sources couldn't confirm them. The three-layer model is a synthesis and hasn't been tested empirically.

A closing thought

AI isn't simply taking jobs. It's changing which tasks people do, and it's doing so first at the point where people have always learned to work. That creates an expertise paradox: deep judgement is worth more than ever, while the routine work that built it is disappearing. Schools can't solve that alone, but they can stop preparing young people to compete with machines at the tasks machines do best. The better preparation is to build durable human capabilities, teach young people to supervise and check AI rather than defer to it, and give them real experience of work in which judgement is formed. That's a different education from one built around content recall, and it's closer to what schools have always said they value.

THINK IT THROUGH

Reflect.

- Where in your setting do students make real judgements with real stakes?
- Which students depend most on school to build the bridge to work, and what do they currently get?
- What would you say to a student worried about the entry-level jobs in their chosen field?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- Our curriculum and assessment reward reasoning and learning to learn, not only recall.
- We teach AI fluency as checking and supervising output, within subjects.
- Students have regular contact with real work through projects, mentoring or placements.
- Students without family networks are given the bridge to employers.
- Caring and vocational pathways are presented as ambitious choices.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

PART SEVEN

Teachers at the Centre

> Based on the paper "Teachers at the Centre: Closing the Teacher AI Gap through Professional Learning, Time and Agency" by Paul Penfold (Proceedings of Future Science Archives, 2026).

Part Six looked at what young people are preparing for. This part looks at the adults who prepare them. It argues that the gap between students' use of AI and teachers' confidence with it is the main bottleneck for everything else schools hope to do.

CHAPTER 32

A mismatch at the heart of school AI

There's a mismatch in how schools are handling AI. Students use it constantly. Teachers, on average, use it less, feel less confident with it, and have had little support to learn it. Yet teachers are the people schools rely on to help young people use AI well.

That matters because the most important AI skill for students isn't operating the tools, which they already do, but judging what the tools produce. Students are confident users and often poor critics: they trust output they should question and rarely check it. The only people in a position to teach that judgement are the adults in the room, and a teacher can't easily teach students to think critically about a tool they avoid themselves (Penfold, 2026) [158].

This part treats the teacher AI gap as the central constraint on AI in schools. Three questions guide it:

- How large is the gap between students' and teachers' use of and confidence with AI, and what causes it?
- What can AI realistically offer teachers, and what should happen to the time it saves?
- What kind of professional learning, and what conditions of agency, close the gap?

Where the material comes from

This part draws on newsletters on teacher development that I wrote for school leaders in 2026 (Penfold, 2026) [104, 134, 110, 126, 158], including summaries of the World Economic Forum's report on education readiness for AI. I read them alongside international survey data on teachers and AI, trial evidence on teacher workload, and research on effective pro-

fessional development. I checked the figures from the newsletters against their original sources and corrected them where they differed.

CHAPTER 33

How big is the gap?

Students are already there

Students' use of AI is near universal in many systems. A 2026 global survey of more than 45,000 students and staff in higher education found that 88% of students used AI, but only 29% believed their instructors were equipped to guide them (Digital Education Council, 2026) [36]. Stanford's 2026 AI Index reported that four in five K–16 students in the United States used generative AI for schoolwork (Stanford HAI, 2026) [175]. Chapter 18 gave these figures in the context of AI literacy.

Teachers are partly there

Teachers' use is lower, more cautious and patchier.

- Use. In the OECD's 2024 Teaching and Learning International Survey (TALIS), about 41% of lower secondary teachers across the OECD reported using AI in their teaching, ranging from 75% in Singapore and the United Arab Emirates to 14% in France (OECD, 2025 [87]; Education International, 2025 [41]). Among those who used it, the most common uses were learning about and summarising a topic (68%) and generating lesson plans or activities (64%).
- Barriers. Among non-users, 75% said they lacked the knowledge and skills to use AI (Education International, 2025) [41].
- Preparation. Only about 38% of teachers reported that AI had been part of their professional learning in the previous year, from about 76% in Singapore to 9% in France (Education International, 2025) [41]. In a US poll, only about four in ten teachers said their school offered AI-related training (NPR, 2026) [84].

- Guidance. Among US teachers whose school provided AI software, only 35% reported a formal policy on teacher use (NPR, 2026) [84]. In a 2025 survey of university faculty, 80% reported a lack of clarity about how AI could be applied in teaching at their institution, and only 4% were fully aware of their institution's guidelines and felt they were comprehensive (Digital Education Council, 2025) [35].
- Belief in benefit. A survey of more than 20,000 teachers in 14 European countries found that 63% used AI to prepare learning materials, but only 16% believed general-purpose AI improves learning outcomes. Confidence in using generative AI had risen from 28% to 42% in a year (Sanoma Learning, 2026) [166].

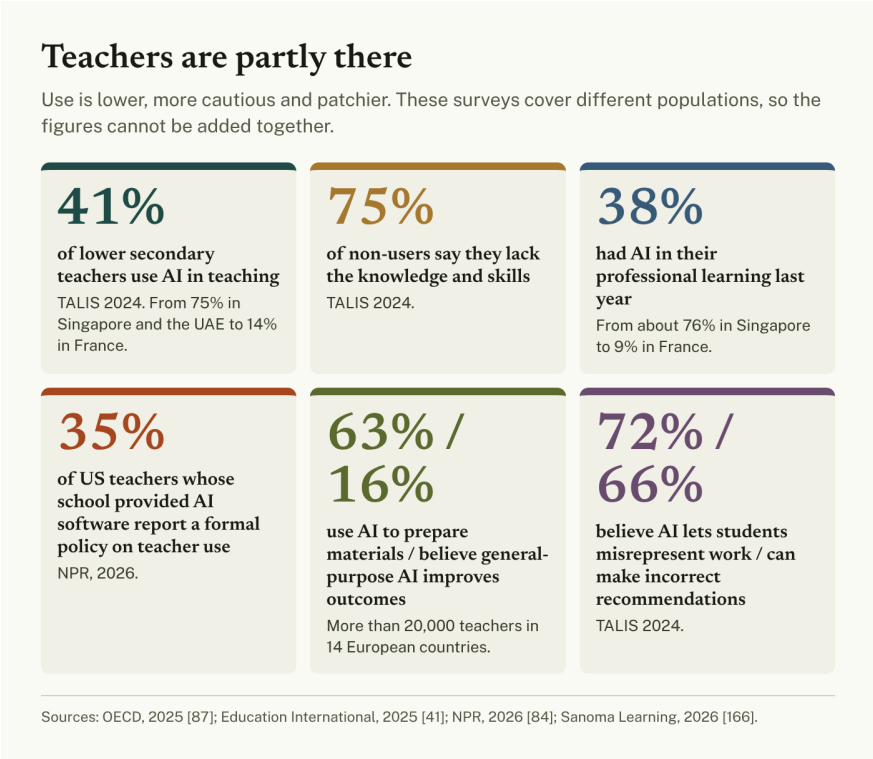


Figure: Teachers' use of and views on AI (TALIS 2024 and other surveys, Chapter 33).

These surveys cover different populations and use different questions, so they can't be added together into a single figure. Together, though, they show a consistent pattern: teachers' use is growing, but it's unevenly supported, often unguided, and accompanied by doubts about whether general-purpose tools help students learn.

Why the gap exists

My newsletters suggest the causes are human more than technical (Penfold, 2026) [158]. Teachers are stretched for time. The tools arrived quickly, with no manual and a great deal of hype. For a year, many were told that AI was mainly a cheating problem to be policed. And there's a less spoken fear of looking less than expert in front of a class that may know the tools better.

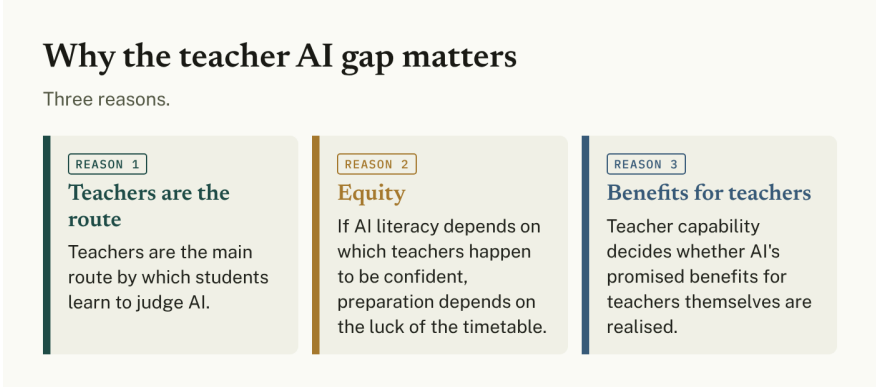


Figure: Why the gap matters (Chapter 33).

Teachers' concerns are also well founded. In TALIS 2024, 72% of teachers believed AI enables students to misrepresent their work, and 66% believed it can make inappropriate or incorrect recommendations (Education International, 2025) [41]. Among university faculty, 83% were concerned about students' ability to critically evaluate AI outputs (Digital Education Council, 2025) [35]. Caution is a reasonable response to a technology that is powerful, unreliable and poorly governed.

Why the gap matters

The gap matters for three reasons. First, teachers are the main route by which students learn to judge AI. Second, it affects equity: if AI literacy depends on which teachers happen to be confident, students' preparation will depend on the luck of their timetable (Penfold, 2026) [158]. Third, teacher capability decides whether AI's promised benefits for teachers themselves are realised.

CHAPTER 34

What AI can offer teachers, and what the time is for

Time

The strongest current evidence for AI in schools concerns teachers' time rather than students' learning. In a trial run by the Education Endowment Foundation in England, 259 Key Stage 3 science teachers in 68 schools were given ChatGPT and a guide to using it. They spent an average of 56.2 minutes a week on lesson planning, compared with 81.5 minutes in a comparison group, a reduction of 31%, with no measurable drop in the quality of the resources produced (Education Endowment Foundation, 2024) [40]. Stanford's AI Hub for Education review similarly found that AI tools can reduce time spent on grading, planning and feedback by as much as 30% without lowering lesson quality (GovTech, 2026) [52]. In the US poll, most teachers who used AI said it saved them time, although 63% said the saving was two hours a week or less (NPR, 2026) [84].

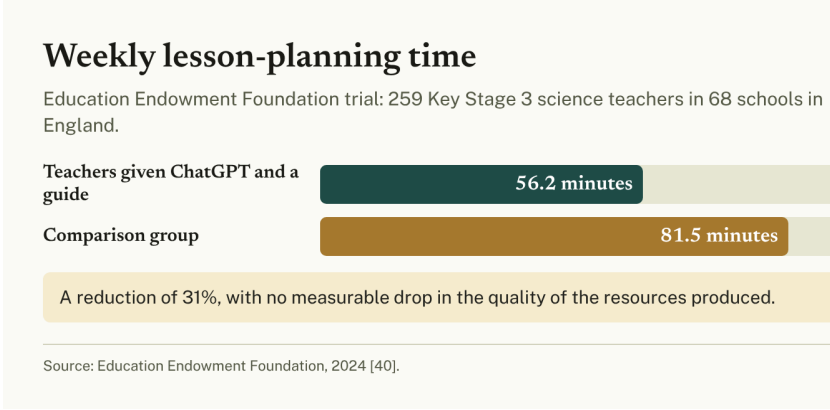


Figure: Lesson-planning time in the EEF trial (Chapter 34).

Scale

The wider context is a global teacher shortage. UNESCO estimates that the world needs 50 million additional teachers by 2030: 44 million for primary and secondary education (UNESCO, 2024) [187] and a further 6 million for pre-primary education, added in UNESCO's October 2026 update (UNESCO and partners, 2026) [209]. The World Economic Forum's readiness report argues that AI's administrative support could help relieve the pressure, citing research suggesting that up to a fifth of teachers' time goes on automatable administrative tasks (World Economic Forum, 2026 [205]; Penfold, 2026 [104]). It also links personalisation to Bloom's (1984) [15] "two sigma problem": one-to-one tutoring produces much better outcomes than group instruction but has never been affordable at scale.



Figure: The teacher shortage and administrative load (Chapter 34).

What the time is for

The WEF report is careful to say that none of these opportunities will be realised automatically (World Economic Forum, 2026) [205]. The conclusion I drew in the newsletters is that the promise of AI isn't that it teaches better than people, but that it might free people to do the parts of teaching that only people can do, such as noticing, encouraging, connecting

and inspiring (Penfold, 2026) [104]. Time saved that gets absorbed into more administration, larger classes or reduced staffing leaves teachers no better off and students worse off. Deciding deliberately how saved time is reinvested is part of any responsible use of AI by teachers.

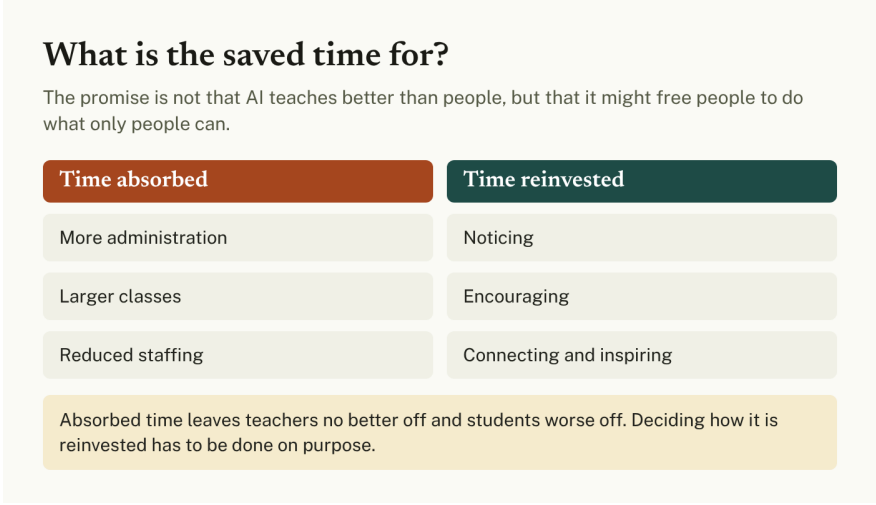


Figure: Time absorbed or time reinvested (Chapter 34).

CHAPTER 35

Professional learning that works

The usual institutional response to the teacher AI gap is a single training day. The research on teacher professional development suggests why that rarely changes practice.

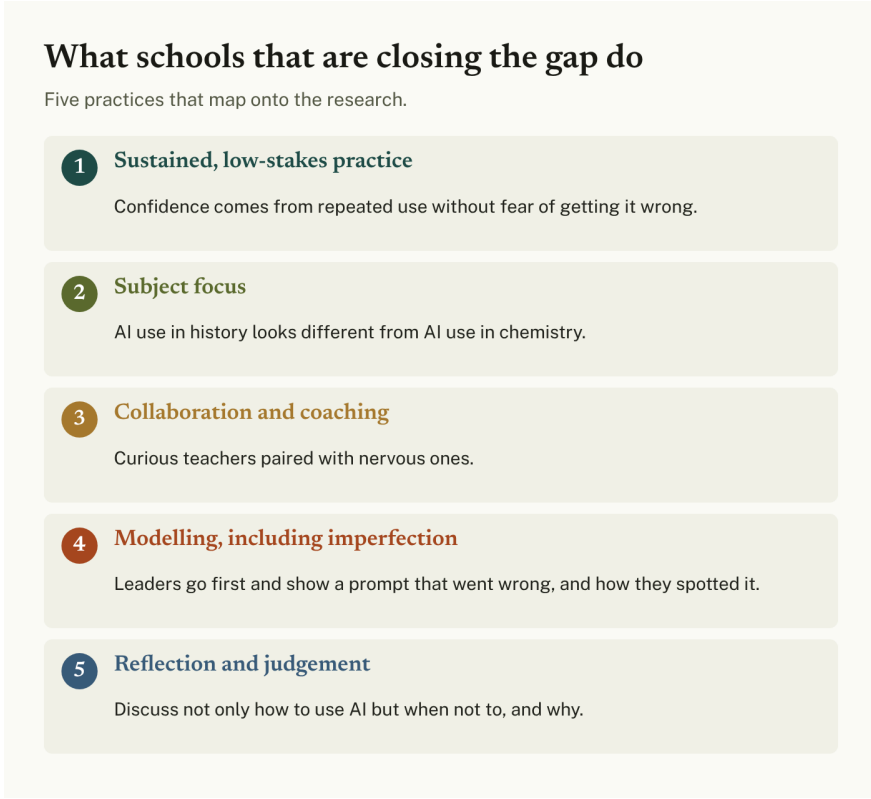


Figure: Five practices of schools closing the gap (Chapter 35).

What the research already shows

Kennedy's (2016) [61] review of professional development programmes found that programmes which help teachers understand and adapt ideas to their own context, rather than prescribing procedures, are more likely to improve teaching. Darling-Hammond, Hyler and Gardner (2017) [30] identified seven features of effective professional development: it is content-focused, involves active learning, supports collaboration, uses models of effective practice, provides coaching and expert support, offers time for feedback and reflection, and is sustained over time. Kraft, Blazar and Hogan's (2018) [65] meta-analysis of causal studies found that coaching had substantial positive effects on teachers' instruction and smaller positive effects on student achievement.



Figure: Seven features of effective professional development (Chapter 35).

What schools that are closing the gap do

The practices described in my newsletters, in schools that are closing the AI gap, map closely onto these findings (Penfold, 2026) [158].

- Sustained, low-stakes practice. Confidence with a new tool comes from repeated use without fear of getting it wrong, not from one dramatic session.
- Subject focus. Staff share what actually worked in their own subject, because AI use in history looks different from AI use in chemistry.

- Collaboration and coaching. Curious teachers are paired with nervous ones, and peer coaching replaces one-off lectures.
- Modelling, including imperfection. Leaders go first and are visibly imperfect. Showing a prompt that produced a poor answer, and how they spotted it, does more for a culture of critical use than any handbook.
- Reflection and judgement. Teachers discuss not only how to use AI but when not to, and why.

CHAPTER 36

Agency, the right to say no, and a capability pathway

Teachers in the room

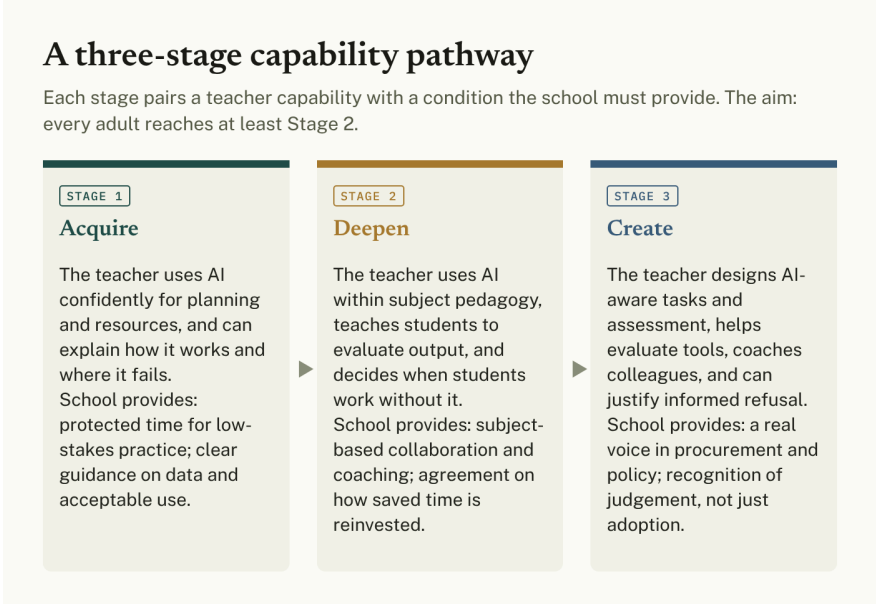


Figure: A three-stage capability pathway for teachers (Chapter 36).

In May 2026, the Council of the European Union adopted conclusions on teachers in the era of AI. They warn against AI that undermines teachers' autonomy and sustainable working conditions, and call for teachers to help design and evaluate the AI tools used in their classrooms (Council of the European Union, 2026 [28]; Penfold, 2026 [134]). The WEF readiness report makes the same argument: AI's value in a classroom depends on the skill and intention of the teacher using it, and top-down rollouts in

which tools are chosen centrally are associated with lower trust and lower uptake (World Economic Forum, 2026 [205]; Penfold, 2026 [110]). Research on teacher agency describes it as something teachers achieve in particular conditions, not a fixed personal trait (Priestley et al., 2015) [160]. The practical implication is that schools create or remove the conditions for teacher agency through their decisions about AI. Chapter 43 returns to the Council conclusions.

Informed refusal as a competence

A further step was proposed at a 2026 workshop of teachers and school leaders convened by UNESCO and the European Union: that teachers need the freedom not to use AI, and that this should be named as a competence (Penfold, 2026) [126]. A teacher who decides to keep a chatbot out of a writing unit because it would undermine the thinking she's trying to build is showing exactly the judgement AI literacy is meant to produce. Most systems measure success by adoption and can't tell her apart from a teacher who has simply been left behind. Professional learning that only ever teaches how to use AI builds a one-way ratchet, in which teachers can move towards more use but can't credibly choose less.

A three-stage pathway

UNESCO's AI competency framework for teachers sets out five dimensions (a human-centred mindset, the ethics of AI, AI foundations and applications, AI pedagogy, and AI for professional development) across three progression levels: acquire, deepen and create (UNESCO, 2024) [186]. Drawing on that framework and on the evidence above, I propose a three-stage pathway for schools. Each stage pairs a capability with a condition the school must provide.

- Stage 1, acquire. The teacher uses AI confidently for personal productivity, such as planning and resource creation, and can explain how it works and where it fails. The school provides protected time for low-stakes practice, and clear guidance on data and acceptable use.

- Stage 2, deepen. The teacher uses AI within subject pedagogy, teaches students to evaluate AI output, and decides when students should work without it. The school provides subject-based collaboration and coaching, and agreement on how saved time is reinvested.
- Stage 3, create. The teacher designs AI-aware tasks and assessment, helps evaluate tools before adoption, coaches colleagues, and can justify informed refusal. The school provides a real voice in procurement and policy, and recognition of judgement, not just adoption.

The aim is that every adult in a school reaches at least Stage 2: able to use AI competently in their own subject, see where it falls short, and help students do the same (Penfold, 2026) [158]. That's a reachable bar. It doesn't require teachers to become computer scientists.

What to do

For school leaders. Treat teacher AI capability as the first stage of any AI strategy, not a parallel workstream. Replace the single training day with sustained, subject-based professional learning that includes coaching. Go first and be visibly imperfect. Decide explicitly how time saved by AI will be reinvested, and protect it from being absorbed by other demands. Map staff confidence so that students' preparation doesn't depend on the luck of their timetable.

For systems and ministries. Fund AI professional learning on the scale of the gap revealed by TALIS, where most non-users cite lack of skills and fewer than four in ten teachers had relevant training. Embed AI in initial teacher education and induction, as the EU Council conclusions recommend. Involve teachers in the evaluation and procurement of AI tools.

For teacher educators and researchers. Evaluate AI professional learning against changes in teachers' practice and students' learning, not just teacher satisfaction or tool adoption. Develop ways of recognising informed refusal as professional judgement.

What this part can't tell you

The survey data come from different populations, years and question wordings and can't be combined. The TALIS figures refer to lower secondary teachers, the Digital Education Council faculty figures to higher education, and the NPR figures to US K–12 teachers. The strongest trial evidence on teacher workload comes from one subject in one country. The links drawn between general research on professional development and professional learning for AI are reasoned rather than tested, and the capability pathway hasn't been evaluated.

A closing thought

Students have already decided that AI is part of their lives. Whether their use becomes wise or lazy, honest or hidden, critical or credulous depends largely on the adults around them, and many of those adults have been given neither the time, the training nor the authority to lead. Closing the teacher AI gap isn't mainly a matter of better tools. It's a matter of sustained professional learning, time that's protected and reinvested in human work, and the agency to decide, including the agency to say no. There's no route to AI-ready students that skips the teachers.

THINK IT THROUGH

Reflect.

- What one change would give teachers a protected hour a week to practise with AI?
- How will time saved by AI be reinvested in your setting, and who decides?
- When would you, as a teacher, decide not to use AI, and how would you explain it?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- Stage 1: teachers have protected time for low-stakes practice and clear guidance on data and acceptable use.
- Stage 2: subject-based collaboration and coaching exist, and we have agreed how saved time is reinvested.
- Stage 3: teachers have a real voice in choosing tools and shaping policy.
- We map staff confidence so student experience does not depend on timetable luck.
- A teacher who decides not to use AI for a good reason is respected, not penalised.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

PART EIGHT

Leading the Redesign

> Based on the paper "Leading the Redesign: Purpose, Ecosystems and Sequencing in School Leadership for AI" by Paul Penfold (Proceedings of Future Science Archives, 2026).

Part Seven put teachers at the centre. This part looks at the people who set the conditions teachers work in. It treats AI in schools as a leadership problem rather than a purchasing one, and sets out an order in which leaders might act.

CHAPTER 37

Buying the tool is the easy part

Buying an AI tool is easy. Building an institution that uses it well is hard, and schools keep skipping straight to the easy part (Penfold, 2026) [124].

The pattern is familiar to anyone who has led change in schools. A framework arrives, well argued and full of good intentions, and then sits unread while timetables, behaviour and budgets carry on as before. Frameworks rarely fail because they're badly designed. They fail because nobody works out how to use them on a Tuesday morning (Penfold, 2026) [112].

This part asks three questions:

- What pressures does AI place on school leadership?
- Why do well-designed changes so often fail to reach classrooms?
- In what order should leaders act?

Where the material comes from, and a declaration

The part draws on newsletters on educational leadership that I wrote for school leaders in 2026 (Penfold, 2026) [124, 112, 145, 152, 141, 95, 131, 140, 148], including summaries of reports from the World Economic Forum, the OECD and UNESCO. I checked figures against primary or authoritative secondary sources, and several were corrected or omitted. I've led international schools, and I'm connected to Brave Generation Academy, an organisation discussed in Chapter 41. Both shape the questions I ask. The Brave Generation Academy connection is a possible conflict of interest in that chapter, and I flag it there.

CHAPTER 38

The speed gap

The World Economic Forum's 2026 education readiness report describes AI as no longer an emerging trend in education but a routine part of how millions of learners find information, complete academic tasks and make sense of the world (World Economic Forum, 2026 [205]; Penfold, 2026 [145]). Adoption has been fast. In RAND's surveys of young people in the United States, the share using AI for homework rose from 48% in May 2025 to 62% in December 2025, while the share who said AI harms critical thinking rose from 54% to 67% (eSchool News, 2026) [42]. Students are adopting a tool they themselves are increasingly worried about.

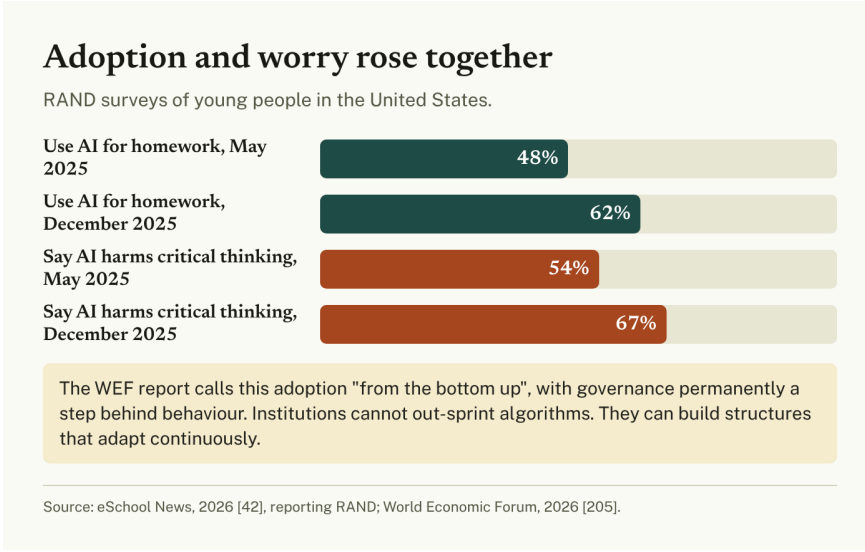


Figure: Adoption and concern among US young people, 2025 (RAND).

The WEF report describes the resulting pattern as adoption "from the bottom up": teachers reacting to students and institutions reacting to

teachers, with governance permanently a step behind behaviour (World Economic Forum, 2026) [205]. Its conclusion is that institutions can't outsprint algorithms and shouldn't try. The realistic response is to build structures that let an institution adapt continuously.

That's a different kind of leadership from drawing up a policy once and defending it. It asks leaders to get comfortable with reviewing their own decisions in public, which is the thread that runs through the rest of this part.

CHAPTER 39

What is school for?

Knowledge as material, not destination

Knowledge: from destination to material

School is not asked to abandon knowledge. People cannot think, create or judge in a vacuum.

Knowledge as destination	Knowledge as material
School as a delivery service for something scarce	What students think with, create with and solve problems with
Knowledge hard to reach	Knowledge in every child's pocket
Purpose: transmission	Purpose widened to skills, attitudes and values, grown in company with others

Figure: Knowledge as destination and knowledge as material (Chapter 39).

The divide is within subjects, not between them

Burning Glass Institute research, as described in Chapter 39.

Shift from execution to direction and editing

Analytical writing, data analysis.

Protect from efficiency pressures

Ethical reasoning, critical reading, empathy.

Keep as foundations

Grammar, calculation. Students cannot direct or evaluate AI without understanding the principles underneath.

Figure: Three kinds of skill within subjects (Chapter 39).

For more than a century, much of school's practical purpose was to transmit knowledge, because knowledge was scarce and hard to reach. That assumption no longer holds. Knowledge is in every child's pocket, and AI can recall, combine and apply it faster than most people (Penfold, 2026) [152]. A school still organised mainly around transmission is running a delivery service for something that no longer needs delivering.

The answer isn't to abandon knowledge. People can't think, create or judge in a vacuum. The shift is that knowledge moves from being the destination to being the material: what students think with, create with and solve problems with. The World Economic Forum's Education 4.0 framing widens the purpose of school to include skills, attitudes and values as well as knowledge (World Economic Forum, 2026) [205]. Those things have to be grown in the company of other people, which puts relationships, character and community closer to the centre of what school is for.

Relational capability

That argument is gaining support from outside education. In July 2026 the Yidan Prize Foundation and Yonsei University convened a conference in Seoul on relational intelligence: the capacity to build trust, bridge differences, resolve conflict and make meaning together (Yidan Prize Foundation, 2026 [208]; Penfold, 2026 [141]). If AI increasingly handles routine cognitive work, the capabilities that remain distinctly human, and increasingly valuable, are relational ones. Schools that treat them as soft extras are reading the evidence backwards. Chapter 31 made a similar point from the labour market.

What students need to be able to judge

Purpose also shapes curriculum within subjects. Research by the Burning Glass Institute argues that the divide isn't between subjects but within them. Some skills, such as analytical writing and data analysis, shift from execution to direction and editing of AI output. Others, such as ethical reasoning, critical reading and empathy, must be protected from efficiency pressures. And foundational skills such as grammar and calculation

remain necessary, because students can't direct or evaluate AI without understanding the principles underneath (Burning Glass Institute, 2026 [20]; Penfold, 2026 [95]).

CHAPTER 40

Why good changes fail: the system around the classroom

The OECD's 2026 paper on an ecosystems approach to curriculum change argues that reform outcomes depend not only on good design but on the setting in which reforms are interpreted and enacted: a web of actors, relationships, routines, resources and infrastructures that is always present, whether or not it's acknowledged (OECD, 2026) [88]. Although the paper doesn't focus on AI, it may be one of the most useful documents for leaders introducing AI, because the same forces that undermine curriculum reform also neutralise new tools.



Figure: Three frictions (Chapter 40).



Figure: Four OECD strategies for change (Chapter 40).

Three frictions

In my reading, three recurring frictions stand out (Penfold, 2026) [131].

- Pedagogical lag. When time and support run out, teachers return to familiar routines.
- Assessment tension. High-stakes examinations that reward recall pull against a curriculum built around competencies, and the curriculum usually loses.
- Structural silos. Teacher education and professional learning evolve separately from what the curriculum asks teachers to do.

Four strategies

The OECD proposes four strategies: non-partisan dialogue among stakeholders to build legitimacy; strengthening the agency of teachers and students; collective impact, in which stakeholders align around a shared agenda and shared measures; and aligning pedagogy, assessment and teacher standards with the intent of the curriculum (OECD, 2026) [88]. Each applies directly to AI. An AI policy that changes what students do but not how they're assessed will lose to the assessment, which is the lesson of Part Five. A tool introduced without professional learning will be absorbed into old routines, which is the lesson of Part Seven.

CHAPTER 41

Sequencing, unbundling and a leadership sequence

Institution before tool

UNESCO-UNEVOC's 2026 practical guide for technical and vocational institutions treats AI as a multistakeholder, institution-wide process, embedded across teaching, curriculum, assessment and operations, with educational purpose, human agency, equity and safety at the centre (UNESCO-UNEVOC, 2026 [190]; Penfold, 2026 [124]). That's a long way from licensing an app. The OECD's Digital Education Outlook 2026 explains why the difference matters: successfully completing tasks with generative AI doesn't necessarily translate into learning (OECD, 2026 [90]). A tool that raises output can erode learning if the institution around it hasn't been redesigned to protect thinking.

Conversation before policy

Before deploying tools, writing policy or booking training, leaders need to ask their community what they actually want AI to do, which is a different question from what it can do (Penfold, 2026) [140]. Teacher readiness, parent trust and student understanding are shaped by dialogue, not by documents issued after decisions have been made.

Red lines before rollout

Clear limits make sustainable adoption possible. New York City's first system-wide guidance, issued in March 2026, used a traffic-light approach. Teachers could use AI for lesson ideas and drafting communications, but AI was barred from assigning grades, determining discipline, and making placement, special education, promotion and graduation decisions

(GovTech, 2026 [51]; Penfold, 2026 [140]). Such a governance floor isn't a restriction on innovation. It's a condition for it, because it settles in advance the questions that would otherwise be fought over after an incident. (Chapter 21 describes what New York did next.)



Figure: New York City's red lines (Chapter 41).

Two decisions, not one

In September 2026, leaders of two of the largest AI developers publicly called for slowing the pace of AI development, and AI-linked shares fell (CNN, 2026 [23]; Penfold, 2026 [141]). That creates a real tension for school leaders, who are often encouraged to use AI to free teachers' time for human connection. In the newsletters I suggested separating two questions that are usually treated as one: what AI should free time for, and how fast the institution should move (Penfold, 2026) [141]. The first is a question of purpose and can be answered now: time saved should go towards the relational work only people can do. The second is a question of risk, and caution from inside the industry is a reason to slow procurement, not to abandon low-stakes uses already doing useful work. One of those in-

dustry leaders is the chief executive of Anthropic, which makes Claude, the AI that co-wrote this book.



Figure: Separating purpose from risk (Chapter 41).

The changing shape of school

Leaders also face a structural shift: families increasingly unbundle schooling, buying the parts that suit their child rather than the whole package (Penfold, 2026) [148]. In several US states, education savings accounts give families a share of public funding to spend on tuition, tutoring, curriculum or therapies (World Economic Forum, 2026) [206]. Elsewhere the same instinct operates without public funding. Tooley's (2009) [179] research documented low-cost private schools serving large numbers of poor families across Africa and Asia. In England, local authorities reported 126,000 children in elective home education in autumn 2025, up from 111,700 a year earlier, with mental health the most common reason given (16%) and philosophical or preferential reasons cited by 12% (Department for Education, 2026) [32].

Hybrid learning hubs, such as those run by Brave Generation Academy, let teenagers study for international qualifications around sport, performance or disruption to conventional schooling. I'm connected to Brave Generation Academy, so weigh what I say about it accordingly.



Figure: The five-step leadership sequence (Chapter 41).

Some of this is less new than it looks. Montessori education has grouped children by their capacity to direct their own learning, rather than strictly by age, since 1907. AI doesn't invent learner-centred models. It gives families and leaders new reasons to notice them. For leaders of conventional schools, the lesson is that families will judge the bundle against its parts. What a school offers that can't be bought separately (community, relationships, belonging and the formation of character) is likely to become its strongest case.

A five-step leadership sequence

Drawing these strands together, I propose a five-step leadership sequence. It's iterative, not linear: each cycle ends with review and returns to purpose.

- Purpose. Decide what the school is uniquely for, and what AI should and shouldn't serve. Could every member of staff give the same answer?
- Conversation. Ask staff, students and families what they want AI to do, and let the answers change decisions.
- Red lines. Set a governance floor: decisions AI won't make, data it won't touch, and who owns incidents.
- Alignment. Line up assessment, professional learning, timetabling and teacher standards with the change, so that the system supports it rather than neutralising it.
- Iteration. As the WEF readiness report recommends, pick one dimension, name the next realistic signal of progress, decide who owns it and set a date to check back (World Economic Forum, 2026 [205]; Penfold, 2026 [112]).

What to do

For school leaders. Start with purpose and conversation, not procurement. Set red lines before rolling out tools. Look for the frictions in your own school, especially where assessment rewards what the curriculum is trying to move beyond. Separate the decision about what AI should free time for from the decision about how fast to move. Be clear about what your school offers that families can't assemble from parts.

For boards and owners. Judge AI strategy by the quality of the human work around it, not by the number of tools adopted. Fund sustained professional learning, not one-off training. Protect time saved by AI from being absorbed into cost-cutting, especially where ownership is focused on growth and margins rather than on teachers and learning.

For system leaders. Align assessment and teacher standards with the curriculum and AI policy you expect schools to implement, as the OECD recommends. Provide governance floors that individual schools can adopt rather than leaving each to invent its own.

What this part can't tell you

This part reflects a practitioner perspective and is based on reports and surveys rather than original empirical research. Much of the evidence comes from the United States and Europe. The leadership sequence is a synthesis and hasn't been evaluated. My connection to Brave Generation Academy is a potential conflict of interest in the section on unbundling. Several figures in the original newsletters were corrected or omitted where they couldn't be confirmed.

A closing thought

AI in schools is often treated as a technology decision. It's more usefully treated as a test of leadership. The tools will keep changing. What lasts is an institution's clarity about what it's for, its relationships with the people it serves, the limits it sets, the coherence of the system around its classrooms, and its habit of honest review. Leaders who sequence the human work first will find that the technology questions become much easier. Those who start with the tools may find that the system around them quietly decides.

THINK IT THROUGH

Reflect.

- In one sentence, what is your school uniquely for?
- Where do your assessment and your curriculum pull in different directions?
- What is one red line you would hold even if every competitor crossed it?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- Purpose: every member of staff could give the same answer about what the school is uniquely for.
- Conversation: staff, students and families have been asked what they want from AI, and the answers changed a decision.
- Red lines: we have written down decisions AI will not make, data it will not touch, and who owns incidents.
- Alignment: assessment, professional learning, timetabling and teacher standards line up with the change.
- Iteration: we have named the next realistic signal of progress, an owner and a date to review.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

PART NINE

Governing AI in Education

> Based on the paper "Governing AI in Education: Uneven Rules, Thin Evidence and Procurement as Policy, 2026" by Paul Penfold (Proceedings of Future Science Archives, 2026).

Part Eight asked how a school leader should act. This part steps back to ask who governs AI in education at all. It follows a busy year of rules and statements, shows how thin the evidence is behind them, and argues that much of the real governance happens in contracts.

CHAPTER 42

Who governs AI in a school?

Ask who governs AI in a school and the answer depends on where you look. At the top sit intergovernmental statements of principle. Below them, a few binding laws and statutory duties. Below those, national and district policy. Then institutional governance: who in a school approves a tool, who handles a concern, who reviews outcomes. And at the bottom, where most of the real decisions are made, procurement: the contract a school signs with a vendor, which settles what data is collected, where it goes, what the tool does with a student's work, and what it costs.

This part looks at how those layers developed during 2026 and how well they connect. My central claim is that they don't yet connect well. Principles have become clearer and more ambitious, but binding duties are patchy and in one major jurisdiction have been delayed. The evidence that would allow schools to choose well remains thin. The result is that much of the governance of AI in education is being done by contract, often without anyone treating it as governance at all.

Three questions guide the part:

- What changed in the governance of AI in education during 2026?
- Where are the gaps between principles, rules, evidence and practice?
- What can schools and systems do now, without waiting for those gaps to close?

Where the material comes from

This is a policy analysis based on documents published or adopted in 2026, read alongside the research and policy literature. Its starting point is a set of newsletters on policy and governance that I wrote for school leaders between May and October 2026 (Penfold, 2026) [134, 101, 111, 123, 115, 146, 137]. They include summaries of UNESCO's collection

The algorithm in the room (UNESCO, 2026) [189] and of the World Economic Forum's report on education readiness for AI. For this book I revisited the primary documents and news sources and checked them, and confirmed dates against official or legal sources where I could. Policy in this area moves quickly, so check anything you rely on against the current position.

CHAPTER 43

A year of rules: what changed in 2026

Stronger international principles

In March, UNESCO, UNICEF and ITU launched a Charter for Public Digital Learning Platforms, setting out principles for publicly governed platforms based on open standards and interoperable architecture, designed in part to avoid lock-in to a single vendor (UNICEF, 2026) [194]. In May, the Council of the European Union adopted conclusions on teachers in the era of AI, the first time the relationship between AI and teaching had been addressed in EU education policy (Council of the European Union, 2026) [28]. The conclusions call for an ethical, safe and human-centred approach, warn of the risk that AI undermines teachers' autonomy and sustainable working conditions, and ask that teachers help design and evaluate the AI tools used in their classrooms (Penfold, 2026) [134]. Chapter 36 looks at what this means for teachers.

In June, the World Economic Forum published an AI readiness framework for education built on enabling foundations, institutional capacities, pedagogical practices and learning experiences (World Economic Forum, 2026) [205]. In September, education ministers meeting at UNESCO's Digital Learning Week adopted a statement calling for education to remain a common good in the age of AI (UNESCO, 2026) [188].

A statutory duty in England

On 7 July 2026, England's Department for Education published Keeping Children Safe in Education 2026, the statutory safeguarding guidance for every school and college in England, in force from 1 September (Department for Education, 2026 [33]; EdTech Innovation Hub, 2026 [39]). The definition of making or sharing nudes and semi-nudes now explicitly includes images that are digitally altered or wholly generated by

AI, so a fabricated image of a pupil is handled through existing child-on-child abuse procedures rather than falling into a policy gap. The guidance directs schools to the department's product safety expectations for generative AI, applies filtering and monitoring duties to AI tools, and sets an expectation that schools be mobile phone-free by default.

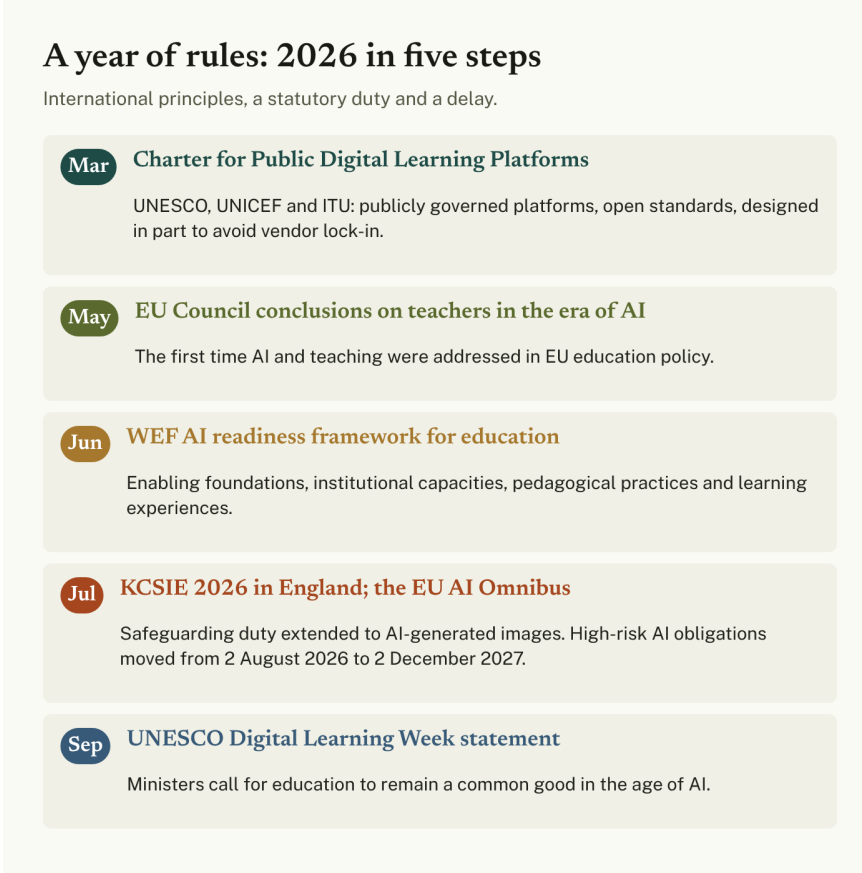


Figure: A year of rules, 2026 (Chapter 43).

This is the first time a national system has placed generative AI inside a binding safeguarding duty rather than alongside it in advice (Penfold, 2026) [101]. Advice can be read and shelved. A statutory duty attaches to a named role.

A delay in Europe

The EU AI Act classifies AI systems used in education for admissions, assessment of learning outcomes, steering of learning and monitoring of students during tests as high-risk (European Union, 2024) [44]. Those systems must meet requirements for risk management, transparency and human oversight, and deployers that are public bodies must carry out a fundamental rights impact assessment.

In July 2026, however, the AI Omnibus, Regulation (EU) 2026/1744, came into force and moved the date from which these high-risk obligations apply from 2 August 2026 to 2 December 2027 (National Law Review, 2026 [79]; Gibson Dunn, 2026 [49]). The prohibited-practice rules and transparency duties for AI systems kept their original dates, so the delay is narrower than headlines suggested. But it covers exactly the uses schools and universities are deploying now: grading, admissions and exam monitoring (Penfold, 2026) [111].

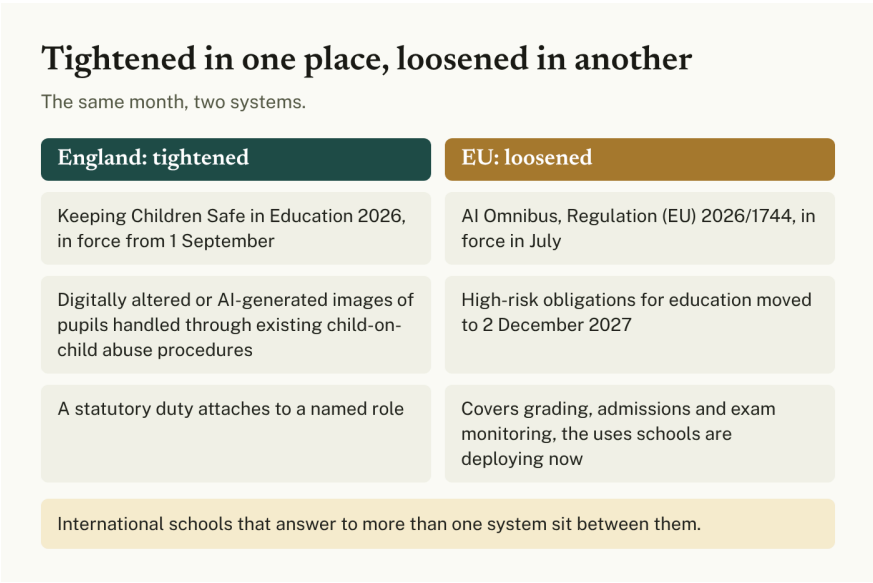


Figure: England tightened, the EU delayed (Chapter 43).

The effect is that one system tightened the safeguarding end of AI in schools in the same month that another loosened the assessment end. School leaders, especially in international schools that answer to more than one system, sit between them.

From strategy to delivery

In the same fortnight, the World Bank opened a Global EdTech Policy Academy on AI in Seoul, with the theme of moving from policy to implementation at scale (World Bank, 2026) [201]. Taken together with the UNESCO statement and the English guidance, three separate institutions signalled that the strategy-writing phase is ending and the delivery phase has begun (Penfold, 2026) [101].

CHAPTER 44

The evidence gap

Governance depends on evidence about what works. Here the picture in 2026 was thin. A review by Stanford's AI Hub for Education examined more than 800 studies on AI in education but found only about 20 high-quality causal studies capable of showing whether AI tools directly change student outcomes (GovTech, 2026) [52]. Its consistent finding was that students often perform better while using AI and struggle to reproduce those results without it, which matches the pattern found in a large field experiment in high school mathematics (Bastani et al., 2025) [11]. The review also found real benefits for teachers, with time spent on grading, planning and feedback reduced by as much as 30% without loss of lesson quality. Its authors described the overall picture as mixed and early.

Practice is running well ahead of this evidence. NPR reported in September 2026 on US districts buying tools first and asking questions later, including a social-emotional chatbot purchased for vulnerable students at an alternative school and a humanoid robot pilot dropped after parents raised privacy concerns (NPR, 2026 [85]; Penfold, 2026 [123]). Teachers are cautious about the general-purpose tools they're using. A survey of more than 20,000 teachers in 14 European countries found that 63% used AI to prepare learning materials, but only 16% believed general-purpose AI improves learning outcomes, and between 75% and 93% said educational AI tools should be designed specifically for education (Sanoma Learning, 2026) [166].

The gap between adoption and evidence is itself a governance problem. Where evidence is thin, the decision about whether a tool is worth using is effectively made by its vendor's marketing and the buyer's hopes.



Figure: The evidence gap (Chapter 44).

CHAPTER 45

Procurement: the hidden layer of governance

My newsletters returned repeatedly to one point: procurement decisions are governance decisions, whether institutions treat them that way or not (Penfold, 2026) [115]. When a school adopts an AI platform, the contract settles questions that would otherwise belong to policy: what student data is collected, where it's stored and for how long, who can see how the tool reaches its decisions, whether the tool's pedagogy fits the school's, what happens when the contract ends, and what it costs over time.

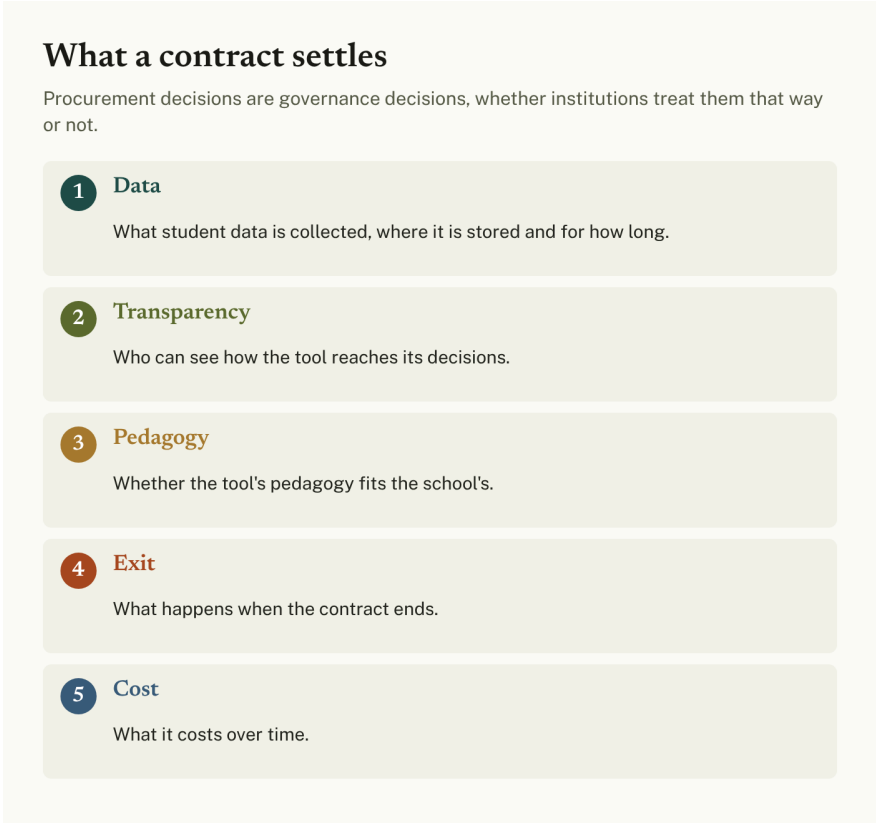


Figure: What an AI contract settles (Chapter 45).

Whose systems are they?

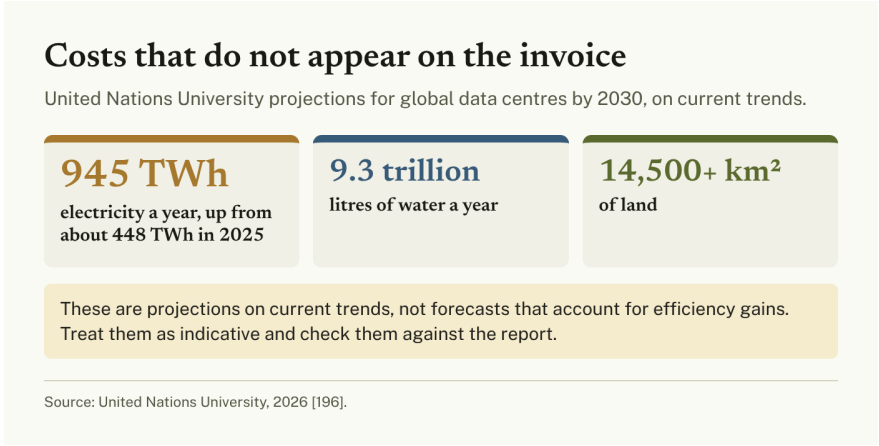


Figure: Projected data-centre resource use, 2030 (United Nations University, 2026).

Contributors to UNESCO's The algorithm in the room make the wider argument. Amiel, Cox and de la Higuera argue that handing educational content to an algorithm hands oversight to a black box, and that open education is the practical route to systems answerable to educators rather than shareholders (Amiel et al., 2026) [5]. Fabo argues that nation-states have traditionally decided curriculum and what counts as reliable knowledge, but that AI companies are taking on more of that role simply by controlling the infrastructure schools depend on (Fabo, 2026) [45]. A ministry can sign a statement about public accountability. Whether it controls the systems the statement is about is another matter.

Who pays, and for what

Appoldt and Gong apply Baumol's cost disease to teaching (Appoldt & Gong, 2026) [8]. Baumol (1967) [12] showed that labour-intensive services become relatively more expensive over time because, unlike manufacturing, their core activity can't easily be automated. Teaching is a classic case: the human interaction is the product. Appoldt and Gong's warning is that when the rising relative cost of education is framed as a problem for AI to solve, the pressure moves towards cuts to staff and privatisation

presented as innovation, rather than towards the harder political choice of funding education properly as a public good (Penfold, 2026) [115].

Costs that don't appear on the invoice

Some costs of AI never appear on a school's budget. A United Nations University report published in June 2026 projected that global data centres could consume 945 terawatt-hours of electricity by 2030, up from about 448 terawatt-hours in 2025, together with about 9.3 trillion litres of water a year and more than 14,500 square kilometres of land (United Nations University, 2026) [196]. These are projections on current trends rather than forecasts that account for efficiency gains, so treat the figures as indicative and check them against the report. They do show the scale of the physical infrastructure that education's AI subscriptions depend on. O'Sullivan and Kelly describe the resulting contradiction for higher education, whose public mission now rests partly on infrastructure with environmental costs it doesn't account for (O'Sullivan & Kelly, 2026 [93]; Penfold, 2026 [146]).

There's an access dimension as well. UNESCO's 2023 Global Education Monitoring Report asked of technology in education: a tool on whose terms? (UNESCO, 2023) [183]. Infrastructure and connectivity remain an equity fault line, and the WEF readiness report warns that institutions with weak infrastructure risk falling further behind, not because their educators are less capable but because the basic plumbing isn't there (World Economic Forum, 2026 [205]; Penfold, 2026 [137]). Part Ten takes up equity directly.

CHAPTER 46

A governance stack, its gaps, and a protocol for schools

Five layers



Figure: The five-layer governance stack (Chapter 46).

Bringing these strands together, the governance of AI in education can be understood as a stack of five layers. Each answers a different question, and each depends on the layers around it.

- 1. International principles. What should AI in education be for? In 2026: the UNESCO ministerial statement, the EU Council conclusions and the Charter for Public Digital Learning Platforms. State: stronger and more specific.
- 2. Binding law and statutory duty. What must institutions do, and who is accountable? In 2026: the EU AI Act high-risk rules and KCSIE 2026 in England. State: uneven, tightened in England and delayed in the EU.
- 3. System policy. How will this system implement AI? Examples: national curricula, district moratoria and procurement frameworks. State: fragmented and fast-changing.
- 4. Institutional governance. Who in this school decides, reviews and answers for AI use? Examples: approval routes, named roles and review cycles. State: often informal or missing.
- 5. Procurement and practice. What exactly does this tool do with our students and their data? Examples: vendor contracts, pilots and classroom use. State: where most real decisions are made.

Three gaps

The binding gap, between layers 1 and 2. Principles have moved faster than law. The EU delay leaves sixteen months in which, for the uses that matter most in education, the main protection between a student and an automated admissions or grading decision is institutional policy (Penfold, 2026) [101].

The capacity gap, at layer 4. The WEF report frames institutional capacity as more than a policy document. It means clear decision-making structures, a route for staff to raise concerns, and governance distributed across leadership, teachers, technical staff and, where appropriate, students and families, rather than concentrated in one person who becomes a bottleneck (World Economic Forum, 2026) [205]. In most schools, a

teacher who finds a promising tool on a Saturday doesn't know who to ask on Monday (Penfold, 2026) [135].

The evidence gap, at layer 5. Procurement decisions are being made with very little evidence about learning outcomes, and with little attention to data terms, pedagogy or environmental cost.



Figure: Three gaps in the governance stack (Chapter 46).

Procurement as governance: six tests

Schools can't close the binding gap, but they can close much of the gap at layers 4 and 5 now. I propose that every significant AI procurement be treated as a governance decision and pass six tests before a contract is signed. The protocol draws on the newsletters' recommendations and on the principles in the EU Council conclusions and the public platforms charter.

- Evidence. Independent evaluation, not case studies, and results on learning measured with the tool switched off (Penfold, 2026) [123].
- Pedagogy. Teachers who will use the tool are involved in choosing and evaluating it, and its design prompts thinking rather than handing over answers (Council of the European Union, 2026) [28].

- Data. Clear terms on what student data is collected, where it's stored, who can access it and what happens to it when the contract ends.
- Accountability. A named role in the school owns the tool, its incidents and its review, and can explain how it reaches decisions about students. Under KCSIE 2026, AI incidents involving pupils are safeguarding incidents and should sit with an existing accountable role (Penfold, 2026) [101].
- Exit. Data can be exported and the school can leave without losing its students' work or records, in line with the public platforms charter's concern with interoperability and lock-in (UNICEF, 2026) [194].
- Hidden costs. The vendor discloses its energy and water footprint, as it would disclose its safeguarding arrangements, and the school considers whether the tool's benefit justifies them (Penfold, 2026) [146].

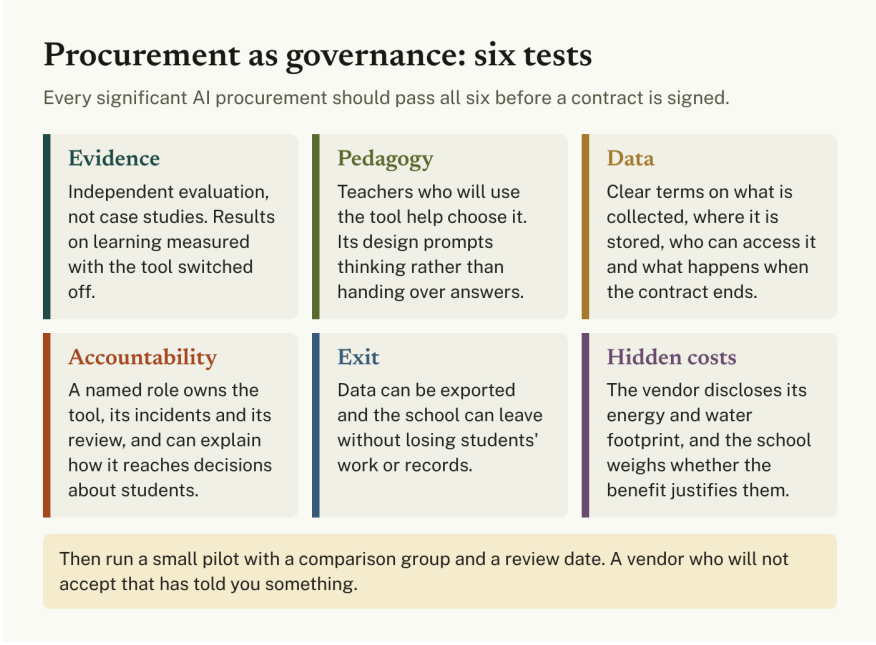


Figure: Six procurement tests (Chapter 46).

A small pilot with a comparison group and a review date should follow. As I put it to school leaders, a vendor who won't accept that has told you something (Penfold, 2026) [123].

What to do

For school leaders. Treat AI incidents involving pupils, including fabricated images, as safeguarding incidents owned by an existing role, and rehearse the response with staff rather than publishing a policy that won't be remembered. Don't treat the EU delay as a reason to shelve readiness work: human oversight, transparency and staff AI literacy were sensible before they were deadlines. Apply the procurement protocol to every significant AI purchase and renewal.

For system leaders and ministries. Close the binding gap for high-risk educational uses rather than relying on institutional goodwill. Fund independent evaluation of the tools schools are buying, including outcomes measured without the tool. Support public, interoperable platforms that reduce lock-in, as the 2026 charter proposes.

For international and multi-jurisdiction schools. Map which binding duties apply to you and adopt the stricter standard where systems diverge. A school that can say plainly how it uses AI in assessment and who checks it is in a stronger position with its community than one waiting to see what regulators eventually enforce (Penfold, 2026) [111].

What this part can't tell you

The analysis covers a single, unusually busy year and focuses on documents from Europe, England, the United States and intergovernmental bodies. It doesn't examine regulation in Asia, Africa or Latin America in comparable depth. Some claims rely on secondary reporting of primary documents, and policy in this area changes quickly. The procurement protocol hasn't been tested systematically and is offered as a practical starting point, not a validated instrument.

A closing thought

The governance of AI in education in 2026 was marked by a widening gap between what international bodies say AI in schools should be for and what binding rules, evidence and institutional capacity can currently deliver. In that gap, contracts are doing the work of policy. Schools don't need to wait for regulators to close it. By treating procurement as governance, giving AI incidents a named owner, and asking for evidence of learning with the tool switched off, they can govern the part of the stack they control. The question for any school leader isn't whether AI is being governed in their institution. It already is, by whoever drafted the last contract.

THINK IT THROUGH

Reflect.

- Who in your school can say, today, which AI tools are in use with pupils?
- If a pupil were harmed by a fabricated image, which existing role would own the response?
- What would you want from a vendor before you believed their claim about learning?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- Evidence: independent evaluation exists, including results measured with the tool switched off.
- Pedagogy: the teachers who will use it helped choose and evaluate it, and it prompts thinking.
- Data: we know what is collected, where it is stored, who can access it and what happens at contract end.
- Accountability: a named role owns the tool, its incidents and its review.
- Exit: data can be exported and we could leave without losing students' work.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

PART TEN

Close the Gap or Widen It?

> Based on the paper "Close the Gap or Widen It? AI, Equity and Learning at the Weakest Link" by Paul Penfold (Proceedings of Future Science Archives, 2026).

Part Nine showed how thinly AI in education is governed. This part asks who benefits. It looks at how AI could narrow or widen educational gaps, and then goes to the place where the stakes are highest: learners in crisis and displacement.

CHAPTER 47

Two plausible futures

Every major change in education carries a hidden danger: that it works well for children who already have advantages and quietly leaves the rest further behind. New technologies, methods and resources tend to reach the well-resourced first, and by the time they reach everyone else the head start has compounded (Penfold, 2026) [97].

AI is the most powerful version of this pattern education has faced. The optimistic case is real: a child with no money for private tuition and no expert at home could, in principle, get individual support that once cost a fortune, in their own language and at their own pace. The pessimistic case is just as real and tends to be the default. Advantaged children get powerful tools and skilled adults to guide them, while disadvantaged children get the cheaper version, nothing at all, or unguided access, which can be worse than none.

Three questions guide this part:

- What determines whether AI narrows or widens educational inequality?
- What does the situation look like for the learners furthest from the frontier, those in crisis and displacement?
- What should schools, systems and organisations ask of every AI decision?

Where the material comes from, and a declaration

This part draws on newsletters I wrote for school leaders in 2026 (Penfold, 2026 [97], 2026 [119]), together with data on connectivity, refugee education, aid and AI for development. I checked figures against primary or authoritative secondary sources. Claims that organisations make about their own programmes are reported as their claims. I work for

Brave Generation Academy (BGA), and its work with displaced learners is mentioned in Chapter 50. That's an interest I want you to know about.

Why access isn't enough

Evidence on learning helps explain why access alone isn't enough. In a large field experiment with high school mathematics students, unrestricted access to a chatbot improved performance during practice but lowered results on an exam taken without it. A tutor-style version designed to prompt thinking avoided that harm (Bastani et al., 2025) [11]. Chapter 10 describes the study in full. A child given a powerful tool with no guidance and no well-designed pedagogy isn't necessarily being helped. The equity question is therefore not only who gets the technology, but who gets well-designed tools and skilled human support to use them well, which advantaged children already have in abundance (Penfold, 2026) [97].



Figure: Two plausible futures for equity (Chapter 47).

CHAPTER 48

Four layers of the AI divide

Access

The most basic divide is access to connectivity and devices. The GSMA reported in 2025 that 96% of the world's population lives within reach of mobile broadband, but 3.1 billion people, 38% of the world, live within coverage and still don't use mobile internet. That usage gap is ten times larger than the coverage gap of about 300 million people (GSMA, 2025) [53]. Connectivity and adoption are different problems. Solving the first doesn't automatically solve the second, which involves affordability, devices, skills, relevance and safety.

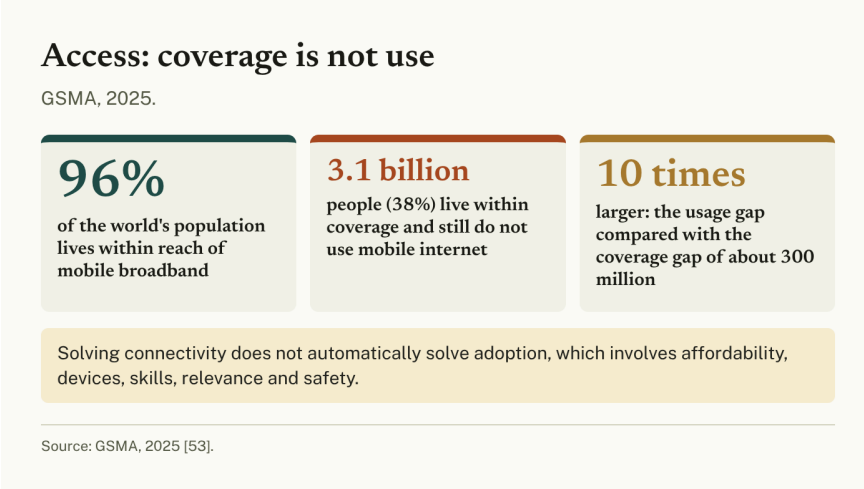


Figure: The usage gap and the coverage gap (GSMA, 2025).

Guidance

The second divide is skilled human guidance. Teachers' own preparation for AI varies widely between systems. In the OECD's 2024 Teaching and Learning International Survey, the share of lower secondary teachers who had AI in their professional learning in the previous year ranged from about 76% in Singapore to 9% in France (Education International, 2025) [41]. Within systems, students in better-resourced schools and homes are more likely to have adults who can help them use AI well. This is why my newsletters argue that equity work should invest most in teaching that helps disadvantaged students use AI wisely, since they're least likely to get that guidance anywhere else (Penfold, 2026) [97]. Part Seven covers the teacher side of this.

Language and culture

The third divide concerns whose language and knowledge AI systems reflect. Research has found that AI writing suggestions nudge writers towards Western styles and reduce cultural nuance (Agarwal et al., 2025) [3]. Contributors to UNESCO's 2026 collection describe AI tools that mishandle local languages and risk overriding local ways of knowing (UNESCO, 2026) [189]. Learners whose languages are poorly represented in training data receive weaker tools, and may also learn that their own ways of expressing themselves are less valued. Chapter 27 raised the same worry about assessment rubrics that reward polish.

Recognition

The fourth divide concerns whether learning is recognised. Learning that can't be certified, or certificates that aren't accepted across borders, leave learners stranded. A young person educated in the camps on the Thai–Myanmar border told Save the Children that their education wasn't recognised, so graduates struggle to find work (Penfold, 2026) [119]. For mobile and displaced learners, recognition may matter as much as access.

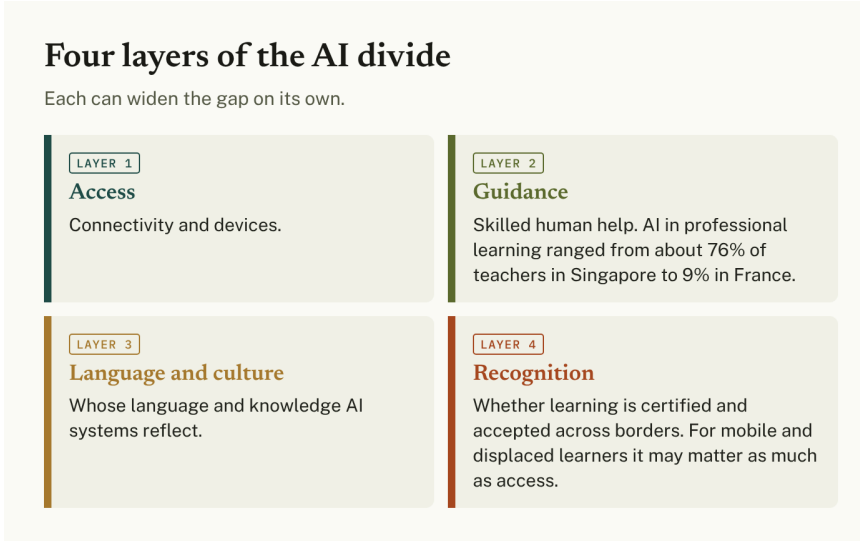


Figure: Four layers of the AI divide (Chapter 48).

CHAPTER 49

The weakest link: learners in crisis

The four layers are hardest for learners in crisis and displacement. UNHCR's 2026 education report estimated that 5.8 million of the world's 13.1 million school-age refugee children, 44%, were out of school in 2025. Gross enrolment was 65% at primary level, 32% at secondary and a record 11% at tertiary level, against a target of 15% by 2030 (UNHCR, 2026) [191].

Each figure hides a particular crisis. In Cox's Bazar, Bangladesh, funding cuts in 2025 led to the closure of more than 6,400 learning centres in the Rohingya camps, putting about 300,000 children at risk of losing their education (Save the Children, 2025) [167]. In Gaza, UNICEF described the education system as on the brink of collapse after two years of war (UNICEF, 2025) [192]. On the Thai–Myanmar border, Save the Children reported in January 2026 that education in the nine refugee camps had reached "breaking point", with aid cuts coinciding with rising enrolment (Save the Children, 2026 [168]; Penfold, 2026 [119]).

The money is moving in the wrong direction. Official development assistance from OECD donors fell by 23.1% in real terms in 2025 to US$174.3 billion, the largest annual fall on record, and humanitarian aid fell by 35.8% (OECD, 2026) [89]. In the same year, world military expenditure rose for the eleventh consecutive year to about US$2.89 trillion (SIPRI, 2026) [169]. These are policy choices, not laws of nature, but they set the conditions in which AI will or won't reach the learners with the least. Please check these figures against the OECD and SIPRI releases before relying on them.

EdTech Hub and IDRC's foresight study of education in conflict and crisis to 2035 ranks a funding collapse as the likeliest scenario, hybrid learning becoming normal as the most urgent, and borderless accredita-

tion for displaced learners as the least likely but among the most consequential (Rahman & Freeman, 2025 [163]; Penfold, 2026 [119]).

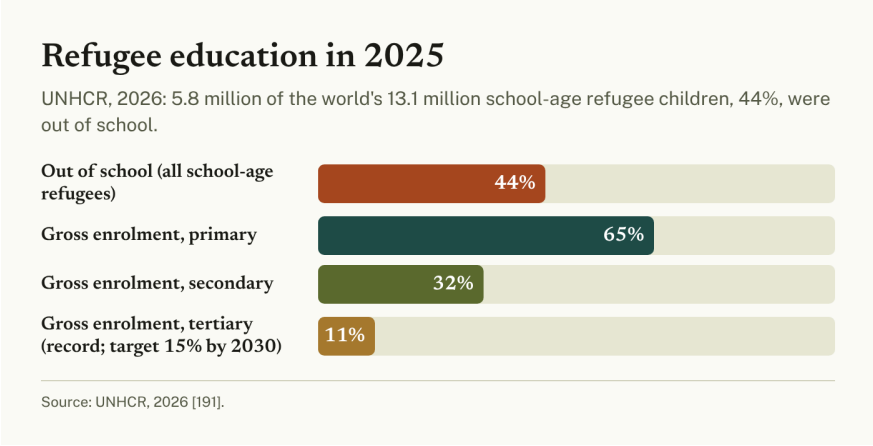


Figure: Refugee education in 2025 (UNHCR, 2026).



Figure: Aid and military spending, 2025 (OECD, 2026; SIPRI, 2026).

CHAPTER 50

What works where power, connectivity and funding fail

If AI is to narrow gaps rather than widen them, it has to work at the weakest link. Several approaches already do, or are trying to.

Offline first

Kolibri, developed by the non-profit Learning Equality, turns a low-cost device into an offline school server that learners reach with their own phones or tablets, and has been used in refugee settings including Kakuma (UNDP, n.d.). In Myanmar, Spring University Myanmar reports using offline Wi-Fi boxes, a learning app and radio lectures to reach students and teachers in conflict areas. It protects the anonymity of teachers and learners because, in a conflict zone, safety is part of the design (Spring University Myanmar, n.d.; Penfold, 2026) [119].

Small AI on basic phones

The World Bank's World Development Report 2026 argues that AI for low- and middle-income countries must often be delivered through voice on basic mobile phones, for people who can't read or can't afford smartphones. It promotes "small AI": affordable models that run on everyday devices (World Bank, 2026) [202].

Supporting the adults

The International Rescue Committee's aprendIA provides low-bandwidth WhatsApp support and short courses for teachers in conflict-affected north-east Nigeria, in English and Hausa (International Rescue Committee, 2026) [59]. Chapter 17 mentioned it as an example of careful

support for carers. In many camps, teachers are refugees themselves. Community teaching isn't a future scenario but a daily reality, and the question is whether it's supported, trained and recognised (Penfold, 2026) [119].



Figure: Design for the weakest link first. Sources: Learning Equality, UNDP Digital X, Spring University Myanmar and EdTech Hub (2025). The descriptions and programme figures come from each project's own pages and are not independent evaluations.

Portable, recognised certificates

The Alsama Project launched the G12++ in 2026, a low-cost, capability-based examination intended to give displaced young people a high-school-equivalent certificate that universities and employers will accept. The organisation reports external benchmarking and early acceptance by universities and employers. These claims are its own and haven't been independently verified (Alsama Project, 2026 [4]; Penfold, 2026 [119]). Brave Generation Academy, with which I'm connected, has worked with partners to train camp teachers in Kakuma as remote teaching assistants (Penfold, 2026) [119].

None of these approaches is automatically safe. EdTech Hub's scenarios warn that private providers may be paid in children's data, and that tools may not fit local languages or cultures (Rahman & Freeman, 2025) [163]. Low cost isn't the same as low risk.

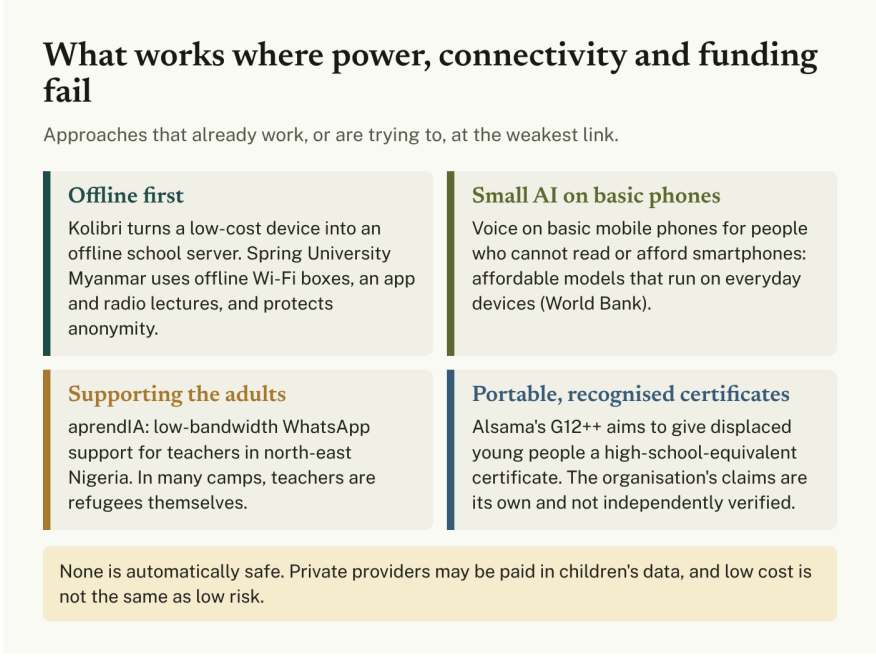


Figure: Four approaches at the weakest link (Chapter 50).

CHAPTER 51

A six-question equity test

Drawing these strands together, I propose six questions that schools, systems and organisations can ask of any AI decision. They extend the simple test I used in the newsletters, "does this narrow the gap or widen it?" (Penfold, 2026) [97], into its component parts.

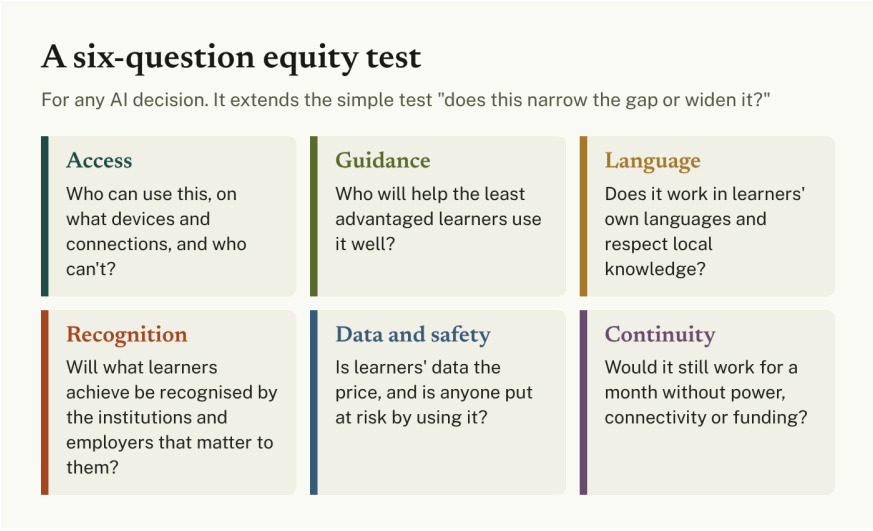


Figure: The six-question equity test as set out in Chapter 51.

- Access. Who can use this, on what devices and connections, and who can't?
- Guidance. Who will help the least advantaged learners use it well?
- Language. Does it work in learners' own languages and respect local knowledge?
- Recognition. Will what learners achieve be recognised by the institutions and employers that matter to them?

- Data and safety. Is learners' data the price, and is anyone put at risk by using it?
- Continuity. Would it still work for a month without power, connectivity or funding?

What to do

For school leaders. Ask of every AI decision whether it narrows or widens the gap. Make sure access to good tools and to skilled guidance doesn't quietly track advantage within your school. Invest most in the students least likely to get support elsewhere.

For systems and donors. Fund the human side, not only the technology: teachers, community educators and the guidance that turns access into learning. Support offline, low-bandwidth and voice-based tools in local languages. Invest in portable, verifiable learner records and recognition across borders.

For organisations working with displaced learners. Plan for lower and less predictable funding. Work offline first. Refuse arrangements in which learners' data is the price. Ask who will recognise a certificate before issuing it. Protect teachers and learners.

What this part can't tell you

This part synthesises data from many sources with different methods and definitions. Several organisational claims about programme reach and recognition couldn't be independently verified. Projections and scenarios aren't forecasts. Much of the evidence on AI and learning comes from well-resourced settings, and its relevance to crisis contexts is uncertain. My involvement with Brave Generation Academy is a potential conflict of interest.

A closing thought

AI won't decide by itself whether it narrows or widens educational gaps. Left to markets and chance, it will probably widen them, because that's how new technologies usually spread. Narrowing them requires deliber-

ate, funded choices about access, guidance, language, recognition, data and continuity, made with the least advantaged learners in mind. The test of AI in education isn't what it does for the best-resourced school, but what it does at the weakest link: in a camp classroom with no power, a teacher who is also a refugee, and a young person whose certificate no one recognises. If it works there, it's likely to work everywhere.

THINK IT THROUGH

Reflect.

- Who is at the weakest link in your own context, and what would help them most?
- Does access to good tools and skilled guidance quietly track advantage in your school?
- What would you ask before accepting a free tool for learners whose data is the price?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- Access: we know who can use this, on what devices and connections, and who cannot.
- Guidance: someone will help the least advantaged learners use it well.
- Language: it works in learners' own languages and respects local knowledge.
- Recognition: what learners achieve will be recognised by the institutions and employers that matter to them.
- Continuity: it would still work for a month without power, connectivity or funding.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

PART ELEVEN

How This Book Was Made

> Based on the paper "From Newsletter to Journal in an Evening: A Human–AI Workflow for Rapid, Open and Verified Publication of Applied Education Research" by Paul Penfold (Proceedings of Future Science Archives, 2026), with an added account of how the book itself was made.

The first ten parts argued for care: with children, with evidence, with contracts, with who gets left out. This last part turns that care on the book itself. It explains how the papers behind these chapters were made, what checking found, what could go wrong with this way of working, and how you can judge it for yourself.

CHAPTER 52

Can good research reach the public in days?

This book began with a simple question: can current, applied research on a fast-moving topic reach the public in days rather than months, without giving up the standards that make research worth reading?

The question matters most where change is fastest. In AI and education, a new model, policy or study can overturn assumptions within weeks. Over 2026 alone, intergovernmental bodies issued new frameworks on AI literacy, England placed AI inside its statutory safeguarding guidance, the European Union delayed its high-risk rules for education, and new evidence appeared on AI tutoring, cognitive offloading and the employment of young workers. A paper written in January and published the following year would miss most of it.

The experiment was to take thematic papers from my newsletter and publish them in a single evening, in an international open-review venue, without peer review at the first stage but with every reference checked, every claim traceable, and the process itself made public. This book is a second step: it takes those papers and turns them into short chapters for readers who won't read a journal article. Eleven papers are behind it, including the process paper this part draws on, which in its published form counts ten.

Why this author, and why this experiment

I've worked in education for six decades, as a teacher, researcher, curriculum and vocational education adviser, school leader and consultant across Asia and beyond. The motivations behind this work are those that have run through that career: a commitment to students' learning, to supporting teachers, to strengthening school leadership and good governance, and, more recently, to understanding how AI is changing learn-

ing, work and young people's futures. Those concerns led to the AI in Education – Deliberately Human newsletter and to a twice-weekly briefing for education leaders, from which the papers are drawn.

The experiment also reflects a longer habit of testing new technologies for learning before they become mainstream. In 2008 and 2009, I built the PolyU Virtual Hotel in the virtual world Second Life at the Hong Kong Polytechnic University, which received a teaching and learning innovation award in 2009. I also published research on teaching and on conducting research inside virtual worlds (Penfold, 2009 [94]; Guillet & Penfold, 2013 [54]). That work explored immersive, networked learning more than a decade before Facebook renamed itself Meta in 2021 and brought the idea of the "metaverse" to mainstream attention. Other work has included development education in Asia, among it refugee education on the Thai–Cambodian border, vocational and hospitality education programmes for governments in the region, and innovation in the schools I've led. Trying new tools early, carefully and in public is a habit. This series applies it to publishing itself.

CHAPTER 53

The speed problem, the money problem and the quality problem

The speed problem

Conventional peer-reviewed publication is slow. A study of publishing delays across disciplines found average times from submission to publication of around a year, with wide variation by field (Björk & Solomon, 2013) [14]. In a slow-moving field that may not matter. In AI and education, a year is long enough for the tools under discussion to be replaced and for the policy landscape to change completely. By the time a paper is published, it risks being out of date.

The economics problem

I also believe that the current academic publishing business is neither ethical nor sustainable, and the evidence supports the concern. Most research is paid for by universities and public funders through researchers' salaries. Researchers then review one another's work largely without payment. One estimate put the time reviewers spent on peer review in 2020 at more than 100 million hours, with the time of US-based reviewers alone worth over US$1.5 billion (Aczel et al., 2021) [2]. Publication itself is concentrated in a small number of commercial publishers, which by 2013 accounted for more than half of all papers in the natural, medical and social sciences (Larivière et al., 2015) [66]. Universities then pay for access through subscriptions, or authors pay article processing charges to make their work open.

The result is that public money pays for research, unpaid academics check it, and institutions pay again to read it. Alternative models that remove charges for both readers and authors are part of the answer.

The quality problem

Speed and openness have their own risks. Generative AI makes it easy to produce text that looks like research. In October 2025, arXiv stopped accepting computer science review and position papers that hadn't passed peer review, after a flood of low-quality submissions it attributed largely to AI (Nature, 2025) [80]. A large-scale study found that the share of scientific papers showing signs of modification by large language models rose quickly after 2022, especially in computer science (Liang et al., 2025) [69]. "Research slop", fluent but shallow or wrong text produced at volume, is a real threat to the scholarly record. Any experiment in faster, AI-assisted publication has to take that threat seriously, and this one tried to.



Figure: The speed, economics and quality problems (Chapter 53).

CHAPTER 54

The venue, and the 81 articles behind it

Future Science Archives

Future Science is a non-profit, open-access and open-source initiative for innovation in academic publishing, hosted at Mohammed VI Polytechnic University in Morocco in partnership with Aix-Marseille Université, the University of Chicago and Columbia University's Center for Contemporary Critical Thought, and led by Raphaël Liogier (Future Science, 2026) [46]. It charges no submission or reader fees and is funded through public and institutional support. It describes itself as being at an early stage of development, moving from an alpha phase towards a wider beta in 2026.

Future Science Archives is one of its initiatives, an Open Review Archive. It welcomes research from all disciplines, posts manuscripts immediately on submission, and invites open commentary and revision in public, like a preprint enriched by community feedback and editorial oversight (Future Science, 2026 [46], 2026 [47]). Its information for authors is explicit about what this means. Submissions aren't filtered in advance for relevance, originality, importance or quality. Publication isn't a certification of the work. Authors remain fully responsible for content, accuracy, methods, citations and ethics, and works may receive open, post-publication peer review. All authors must follow a code of conduct requiring academic integrity, proper citation, transparency, authorship and conflict-of-interest disclosure (Future Science, 2026) [47].

Why it fitted, and where it doesn't

The venue fitted the experiment for three reasons. It publishes immediately, which made a single-evening test possible. It's free to authors and readers, which matched my view of how publicly funded knowledge

should circulate. And it makes evaluation public after publication, which places the responsibility for quality squarely on the authors and invites others to test it.

It also has real limitations, which you should weigh. Its publications aren't indexed in commercial citation databases and carry no impact factor, which matters for early-career researchers whose careers depend on such measures. Experiments with automated AI research and review risk hallucination, bias and a flood of unvetted synthetic literature if safeguards fail. Its long-term sustainability depends on continued public and institutional funding and goodwill. Several features, including automated DOI assignment and AI-assisted reviewing, were still in development at the time the process paper was written, and the paper's own DOI field read "TBD". For now it works better as a research preview than as a primary outlet for mainstream work.

The series didn't begin there. The first article was prepared for a different Future Science initiative, Mirror, which publishes research by AI systems about AI systems. An education review doesn't fit that scope, so I redirected the work to the multidisciplinary archive. That decision is an example of why judgement about venue can't be delegated.

The source corpus

The papers draw on 81 long-form articles I published, mostly during 2026, and archived on my website at paulpenfold.com/blog (Penfold, 2026) [117], alongside a twice-weekly briefing for education leaders. The website tags each article with one or more of eight themes. Because most articles carry two or more tags, the theme counts total 164 rather than 81. According to the process paper, the first ten papers cite 60 of these articles directly, together with an unpublished manuscript and a newsletter report I shared.

The newsletters were written with AI assistance for research, drafting and graphics, and each carried a disclosure to that effect. They draw on reports from UNESCO, the OECD, the World Economic Forum, the European Union, national governments and research institutions, and on my own experience. Their purpose was timely interpretation for busy

school leaders, not academic publication. That difference turned out to matter, as the next chapter shows.

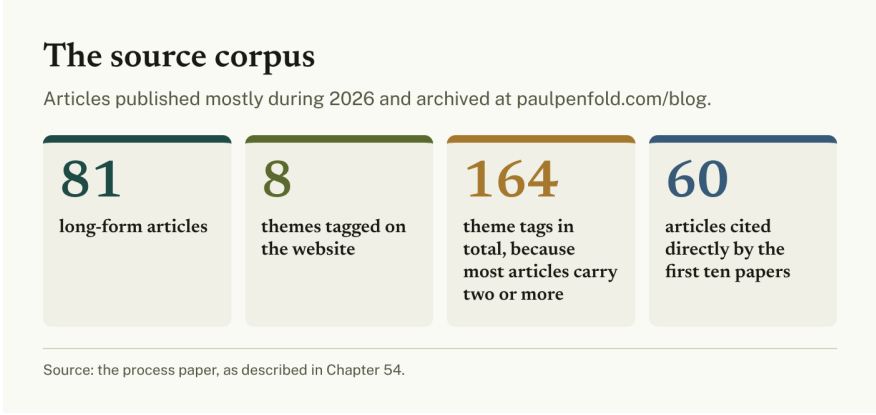


Figure: The source corpus behind the papers (Chapter 54).

CHAPTER 55

The workflow, and what checking found

Who did what

The work was a joint project. I did the thinking that an AI couldn't do: choosing the themes, writing and curating the newsletters over many months, doing the initial research and reading, deciding the venue, setting the standards, reviewing every paper, approving or changing it, and taking responsibility for it. The AI agent (Claude, made by Anthropic) did much of the heavy lifting: gathering the articles, tracing sources, checking references, synthesising each theme into an argument, drafting the text, drawing the figures, and submitting the approved papers through the archive's interface.

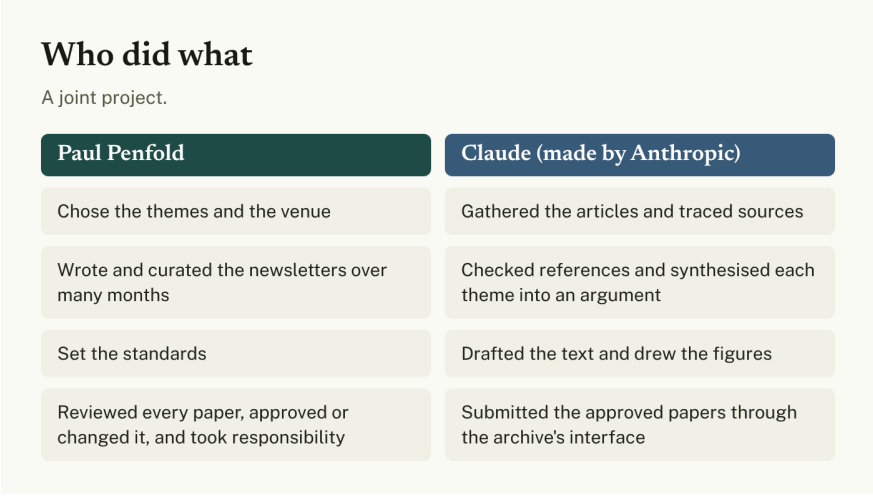


Figure: Who did what (Chapter 55).

The stages

- Write and curate. Over 2025 and 2026, I published 81 long-form articles and regular briefings.
- Choose themes and venue. I selected the themes and the venue.
- Gather. The agent collected the articles by theme from my website and retrieved their full text from their original publication pages.
- Synthesise. For each theme, the agent proposed a distinct argument and structure, avoiding overlap with the other papers, and I steered it.
- Verify. The agent traced each claim to its primary source where possible: checking DOIs against Crossref, reading official documents and press releases, and confirming figures. Claims that couldn't be verified were corrected, attributed or removed.
- Draft. The agent wrote each paper in a conventional academic structure, in plain British English, with the newsletters cited as the source series and a disclosure of how it was made.
- Visualise. The agent drew original figures and a cover image for each paper, inspected every figure and redrew any with faults.
- Review and approve. I reviewed each paper and its figures. Nothing was submitted without my explicit approval, and several papers were changed at this stage.
- Submit and check. The agent submitted each approved paper and then confirmed through the archive's public interface that the paper and every image had published correctly.
- Revise in public. When errors or improvements were found after publication, the agent submitted a revised manuscript linked to the original, so the record of change stays visible.

The process wasn't linear. Three loops ran throughout: a verification loop, in which each claim was checked and the draft revised; a figure loop, in which each figure was rendered, inspected and redrawn until it was clean; and an approval loop, in which I reviewed, asked for changes and approved.



Figure: The ten stages of the workflow (Chapter 55).

What verification found

The most important finding of the experiment is how much the verification step mattered. The newsletters had been written with care, with AI assistance and with sources listed. Yet tracing their claims to primary sources found more than two dozen problems, of the kinds that a fast, uncritical conversion of newsletters into papers would have published.

- A retracted study. A widely cited 2025 meta-analysis reporting large benefits of ChatGPT for learning had been retracted. It was removed, and the retraction was noted as a caution in its own right.
- Misreported findings. One newsletter reversed the findings of a well-known field experiment, attributing a performance loss to students

who didn't use AI rather than to those who did. Another credited a finding from one study to a different study.

- Figures that differed from their sources. Examples included a survey reported as showing that 6% of teachers found AI guidance clear, where the source said 4%, and an employment estimate from an older version of a study that had since been updated.
- Claims that couldn't be verified. These included survey percentages that couldn't be traced to any primary source, a named legislative bill, a government funding figure, a quotation attributed to a named expert, and an article that couldn't be found. All were removed.
- Sources that couldn't be confirmed. A cited book chapter and a named assessment framework couldn't be found in any public record, and were removed or described without attribution.
- Misattribution. A programme run by one humanitarian organisation had been attributed to another.

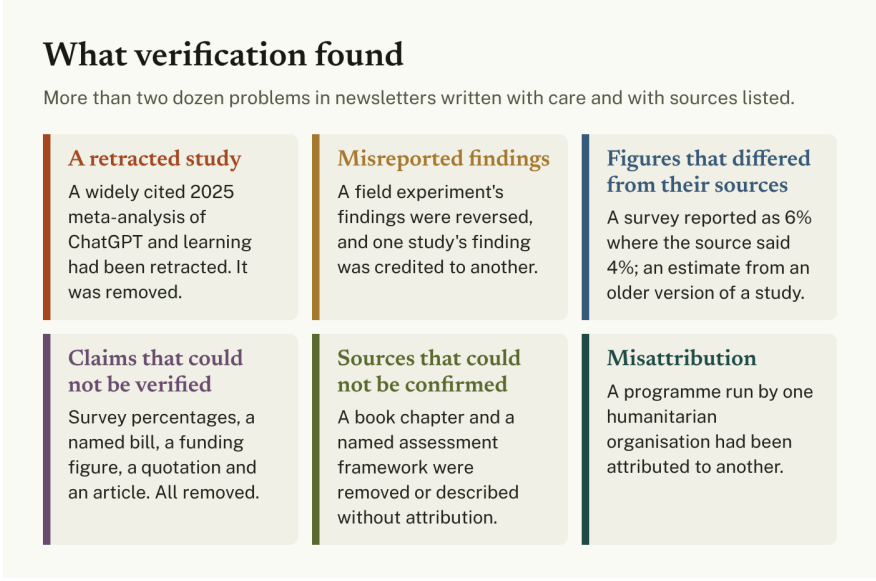


Figure: What verification found (Chapter 55).

Two lessons follow. First, AI-assisted writing can carry errors confidently, and errors persist when text is reused. The newsletters were themselves partly AI-assisted, and some of their errors were the kind AI systems are known to make. Second, AI can also be an effective checker when it's instructed to trace claims to primary sources and to say plainly what it can't confirm. The same technology that creates the risk of research slop can, used deliberately, reduce it. What it can't do is decide what matters, which sources deserve trust, or when a claim should be dropped rather than softened. Those were judgements I made or approved.

The timeline

The first article in the series was published on 25 October 2026. The main series was produced on the evening of 29 October 2026, Thailand time. According to the process paper, the first of the thematic papers was published at 20:40 and the ninth at 22:18 (UTC+7), with revised manuscripts for earlier papers interleaved as figures and corrections were added. Allowing for the earlier preparatory work, the experiment showed that verified, illustrated papers could move from an existing body of practitioner writing to public, citable availability in a single evening.

That speed rested on a large amount of earlier work. The newsletters represent months of reading, thinking, writing and engagement with readers. The evening was the final step, not the whole of the effort.

CHAPTER 56

Risks, conditions and how this book was made

What this approach is, and isn't

These papers are syntheses and arguments grounded in published evidence and practitioner experience. They aren't new empirical studies. At the time of writing they hadn't been peer reviewed, because the venue's model is to publish first and review openly afterwards. Treat them as working papers that invite challenge, which is what the venue is designed for. The same goes for this book.

The risks

The risks are real. If this approach became common without care, it could add to the flood of fluent, low-value text that already threatens the scholarly record. Academic rigour could suffer if verification is skipped. And researchers could become editors of machine-written text rather than authors of their own thinking, losing the slow, effortful work through which understanding is built. That's the same risk several parts of this book describe for students.

Four conditions

The experiment suggests that these risks can be managed when four conditions hold.

- The ideas come from sustained human work, here a long-running newsletter grounded in my reading and experience.
- Every claim is traced to a source, and unverifiable claims are removed rather than softened.

- A human with relevant expertise reviews and approves each paper and remains responsible for it.
- The process is disclosed in full, including the role of AI, the author's interests and the errors found.

Under those conditions, AI-assisted rapid publication isn't a shortcut around thinking. It's a way of getting the results of thinking into the public domain while it's still useful.

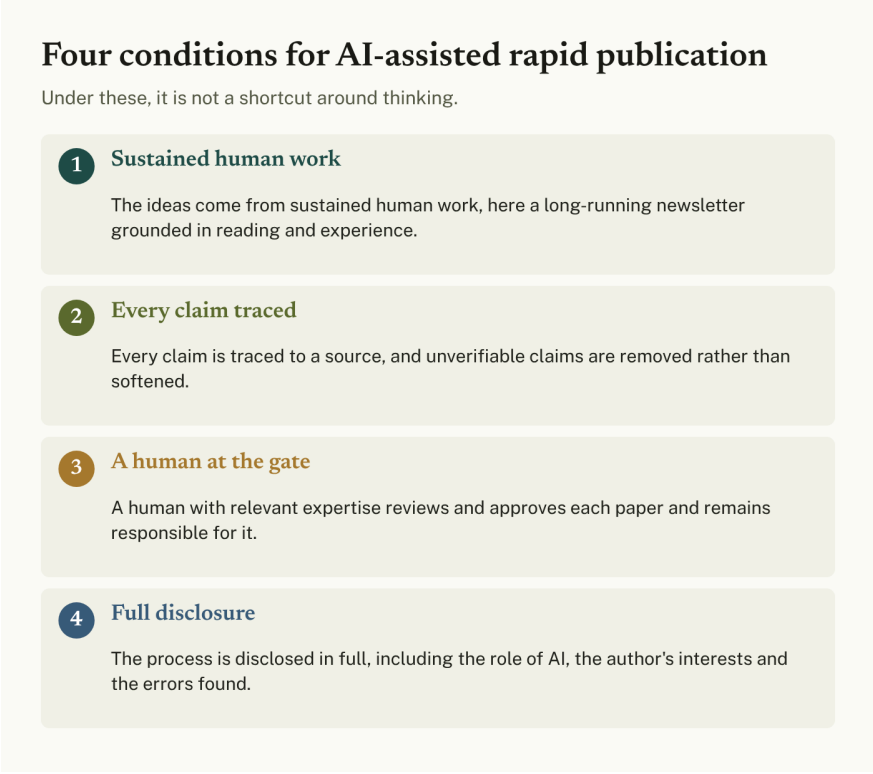


Figure: Four conditions (Chapter 56).

How this book was made

This book is a further step, and it has its own limits, so here is the plain account. The structure of eleven parts and 56 short chapters was agreed

first. Claude then drafted each part from the matching paper, keeping to the paper's evidence and wording where the evidence was concerned, and adding cross-references so that studies described in one part point to the part where they're treated in full. The sources were taken from the papers' reference lists, not rebuilt from memory, and merged into one numbered Bibliography at the back. For Part Two, six more newsletters from my blog were added as my own work. I've reviewed and approved the text, and I'm responsible for it. The chapters are shorter than the papers they come from, because readers who aren't specialists tend to prefer short chapters. Some detail from the papers, including the figures, was left out to keep them short.

Because the book was drafted from the papers, any error in a paper will be in the book too, unless it was caught. Readers who rely on a figure should go back to the cited source. The aim is to correct in the open, as with the papers.

Conflicts of interest

The AI that co-wrote this book is Claude, made by Anthropic, an AI developer whose products and leaders are discussed in several parts. I'm connected to Brave Generation Academy, which is discussed in Parts Eight and Ten. These interests are declared where they arise.

A replicable workflow

For others who want to try something similar, the experiment suggests a simple sequence.

- Start from sustained work of your own. A body of writing, data or practice that reflects real thinking over time.
- Choose an open venue that fits. Check its scope, its review model and its requirements before you start.
- Synthesise by theme, not by article. Build each paper around an argument, and avoid overlap between papers.

- Verify everything. Trace each claim to its primary source, check DOIs, and remove what can't be confirmed.
- Make the evidence visible. Use original figures, and label claims by the strength of their evidence.
- Keep a human at the gate. Nothing is published without the responsible author's review and approval.
- Disclose fully. State the role of AI, the authors' interests and the limitations.
- Correct in public. Use the venue's revision mechanism, so the record of change stays visible.



Figure: A replicable workflow (Chapter 56).

The papers behind this book

All are published in Future Science Archives under my name, with Claude disclosed as co-writer. The process paper gives the web address of each of the first ten.

- Learn first, then use AI: Testing an expert consensus on generative AI and learning against the empirical evidence (Part Two).

- What counts as AI literacy? A comparative analysis of international frameworks and system-level policies, 2024–2026 (Part Four).
- When the artefact is no longer evidence: Assessment validity and academic integrity in the age of generative AI (Part Five).
- Governing AI in education: Uneven rules, thin evidence and procurement as policy, 2026 (Part Nine).
- Teachers at the centre: Closing the teacher AI gap through professional learning, time and agency (Part Seven).
- Growing up with machines that talk back: Wellbeing, attachment and the developing mind in AI-mediated childhood (Part Three).
- The broken bottom rung: Entry-level work, durable capability and what schools should prepare young people for (Part Six).
- Leading the redesign: Purpose, ecosystems and sequencing in school leadership for AI (Part Eight).
- Close the gap or widen it? AI, equity and learning at the weakest link (Part Ten).
- Educational futures to 2035 and 2045: Scenarios, switches and the choices to make now (Part One).
- From newsletter to journal in an evening: A human–AI workflow for rapid, open and verified publication of applied education research (this part).

What this part can't tell you

The workflow described here comes from one author, one body of work and one evening, and hasn't been independently evaluated. The counts of papers and articles come from the process paper and need updating where the eleventh paper changes them. Some details of the venue, such as DOI assignment and indexing, were changing at the time of writing.

A closing thought

The question that began this work was whether current, applied research could reach the public in days rather than months without abandoning rigour. For papers on AI in education, the answer was yes, with conditions. The papers were published in a single evening, in a venue that charges no one, with verified references and full disclosure. They were possible because of months of human thinking that preceded them, and they were trustworthy only to the extent that every claim was checked and a responsible human approved every word. The most striking lesson is that verification found more than two dozen problems in material that had already been carefully written. Speed is easy. Speed with care is the harder and more valuable thing, and it's what this experiment set out to test. Whether it succeeded is now for you to judge, in the open.

THINK IT THROUGH

Reflect.

- Where in your own work does AI help you write faster, and where might it hide errors?
- What would a reader need to see to trust a text you made with AI?
- Who is the human who is responsible for what you publish?

Self-check. How far is each of these true where you work: not yet, partly, mostly or fully?

- The ideas come from sustained human work, not generated from nothing.
- Every claim is traced to a source, and unverifiable claims are removed, not softened.
- A human with relevant expertise reviews and approves it and remains responsible.
- The process is disclosed in full, including the role of AI, the author's interests and the errors found.
- Errors found after publication are corrected in public, so the record of change stays visible.

The online edition has an interactive coach for each part, with a scenario, scored feedback and questions answered from the text.

References

[1] 13abc. (2026, April 15). Private colleges face closure risk as new report estimates 25% will close in next 10 years. https://www.13abc.com/2026/04/15/private-colleges-face-closure-risk-new-report-estimates-25-will-close-next-10-years/ Cited in: Ch 2.

[2] Aczel, B., Szaszi, B., & Holcombe, A. O. (2021). A billion-dollar donation: Estimating the cost of researchers' time spent on peer review. Research Integrity and Peer Review, 6, 14. https://doi.org/10.1186/s41073-021-00118-2 Cited in: Ch 53.

[3] Agarwal, D., Naaman, M., & Vashistha, A. (2025). AI suggestions homogenize writing toward Western styles and diminish cultural nuances. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (pp. 1–21). ACM. https://doi.org/10.1145/3706598.3713564 Cited in: Ch 27, Ch 48.

[4] Alsama Project. (2026). G12++. https://alsamaproject.com/g12 Cited in: Ch 50.

[5] Amiel, T., Cox, G., & de la Higuera, C. (2026). Open education principles: Resisting the metrics of AI black boxes. In UNESCO, The algorithm in the room (pp. 109–115). UNESCO. https://doi.org/10.54675/PMRO5023 Cited in: Ch 45.

[6] Anthropic. (2025, April 8). Anthropic education report: How university students use Claude. https://www.anthropic.com/news/anthropic-education-report-how-university-students-use-claude Cited in: Ch 24.

[7] Anthropic. (2026). Claude's constitution. https://www.anthropic.com/constitution Cited in: Ch 16.

[8] Appoldt, G., & Gong, Z. (2026). Baumol's cost disease: Long-term economic implications where machines cannot replace humans. In UNESCO, The algorithm in the room (pp. 116–120). UNESCO. https://doi.org/10.54675/PMRO5023 Cited in: Ch 45.

[9] Associated Press. (2026, June). Sweden set to ban mobile phones in schools, joining trend of shelving screens for students. WTOP. https://wtop.com/lifestyle/2026/06/sweden-set-to-ban-mobile-phones-in-schools-joining-trend-of-shelving-screens-for-students/ Cited in: Ch 17.

[10] Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. Quarterly Journal of Economics, 118(4), 1279–1333. https://doi.org/10.1162/003355303322552801 Cited in: Ch 28.

[11] Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. Proceedings of the National Academy of Sciences, 122(26).

https://doi.org/10.1073/pnas.2422633122 Cited in: Ch 2, Ch 10, Ch 15, Ch 26, Ch 44, Ch 47.

[12] Baumol, W. J. (1967). Macroeconomics of unbalanced growth: The anatomy of urban crisis. American Economic Review, 57(3), 415–426. https://www.jstor.org/stable/1812111 Cited in: Ch 45.

[13] Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe & A. Shimamura (Eds.), Metacognition: Knowing about knowing (pp. 185–205). MIT Press. Cited in: Ch 9.

[14] Björk, B.-C., & Solomon, D. (2013). The publishing delay in scholarly peer-reviewed journals. Journal of Informetrics, 7(4), 914–923. https://doi.org/10.1016/j.joi.2013.09.001 Cited in: Ch 53.

[15] 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), 4–16. https://doi.org/10.3102/0013189X013006004 Cited in: Ch 34.

[16] Bloomberg. (2026, June 29). China expands AI education across all schools in Xi's technology push. https://www.bloomberg.com/news/articles/2026-06-29/china-pushes-all-school-levels-to-teach-ai-in-xi-s-tech-drive Cited in: Ch 20, Ch 21.

[17] Bonnesen, K., Krygsman, A., Hobson, S., Korczak, D., & Vaillancourt, T. (2026). Affective generative artificial intelligence use and youth mental health. JAMA Pediatrics. https://jamanetwork.com/journals/jamapediatrics/article-abstract/2853453 Cited in: Ch 14.

[18] Brookings Institution. (2026, January 14). A new direction for students in an AI world: Prosper, prepare, protect. Center for Universal Education, Global Task Force on AI in Education. https://www.brookings.edu/articles/a-new-direction-for-students-in-an-ai-world-prosper-prepare-protect/ Cited in: Ch 3.

[19] Brynjolfsson, E., Chandar, B., & Chen, R. (2026). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence (August 2026 update). Stanford Digital Economy Lab. https://digitaleconomy.stanford.edu/news/canariesaug26/ Cited in: Ch 2, Ch 29.

[20] Burning Glass Institute. (2026). Which skills matter now? https://www.burningglassinstitute.org/research/whichskillsmatternow Cited in: Ch 39.

[21] Centers for Disease Control and Prevention. (2024). Youth Risk Behavior Survey, United States, 2023. Morbidity and Mortality Weekly Report Supplements, 73(4). https://www.cdc.gov/mmwr/volumes/73/su/su7304a9.htm Cited in: Ch 14.

[22] Chase, C. C., Chin, D. B., Oppezzo, M. A., & Schwartz, D. L. (2009). Teachable agents and the protégé effect: Increasing the effort towards learning. Journal of Science Education and Technology, 18(4), 334–352. https://doi.org/10.1007/s10956-009-9180-4 Cited in: Ch 27.

[23] CNN. (2026, September 14). AI stocks slide after top industry CEOs call for slowdown of technology's development. https://us.cnn.com/2026/09/14/business/ai-stocks-slide-slowdown-development-amodei-altman-intl Cited in: Ch 41.

[24] CNN. (2026, September 2). Nation's largest school district bans AI in the classroom through eighth grade. https://www.cnn.com/2026/09/02/tech/new-york-city-classroom-ai-ban Cited in: Ch 21.

[25] CNN. (2026, August 2). The AI cheating scandal rocking Mexico's most prestigious university. https://www.cnn.com/2026/08/02/americas/mexican-university-unam-ai-cheating-scandal-intl-latam Cited in: Ch 26.

[26] Common Sense Media. (2025). Talk, trust, and trade-offs: How and why teens use AI companions. https://www.commonsensemedia.org/research/talk-trust-and-trade-offs-how-and-why-teens-use-ai-companions Cited in: Ch 14.

[27] Council of Europe. (2026). Recommendation CM/Rec(2026)12 of the Committee of Ministers to member States on artificial intelligence literacy (adopted 2 September 2026). https://search.coe.int/cm/eng?i=09125948802cde8e Cited in: Ch 19.

[28] Council of the European Union. (2026). Council conclusions on teachers in the era of artificial intelligence (Document 9003/26). https://data.consilium.europa.eu/doc/document/ST-9003-2026-INIT/en/pdf Cited in: Ch 36, Ch 43, Ch 46.

[29] Covington & Burling. (2026, May). Senate Judiciary Committee advances GUARD Act regulating minor use of AI. Global Policy Watch. https://www.globalpolicywatch.com/2026/05/senate-judiciary-committee-advances-guard-act-regulating-minor-use-of-ai/ Cited in: Ch 17.

[30] Darling-Hammond, L., Hyler, M. E., & Gardner, M. (2017). Effective teacher professional development. Learning Policy Institute. https://learningpolicyinstitute.org/product/effective-teacher-professional-development-report Cited in: Ch 35.

[31] Dawson, P., Bearman, M., Dollinger, M., & Boud, D. (2024). Validity matters more than cheating. Assessment & Evaluation in Higher Education, 49(7), 1005–1016. https://doi.org/10.1080/02602938.2024.2386662 Cited in: Ch 23, Ch 26.

[32] Department for Education. (2026). Elective home education, autumn term 2025/26. https://explore-education-statistics.service.gov.uk/find-statistics/elective-home-education/2025-26-autumn-term Cited in: Ch 41.

[33] Department for Education. (2026). Keeping children safe in education 2026. https://www.gov.uk/government/publications/keeping-children-safe-in-education--2 Cited in: Ch 43.

[34] Digital Education Council. (2025). Digital Education Council AI literacy framework. https://www.digitaleducationcouncil.com/post/digital-education-council-ai-literacy-framework Cited in: Ch 19.

[35] Digital Education Council. (2025). What faculty want: Key results from the Global AI Faculty Survey 2025. https://www.digitaleducationcouncil.com/post/what-faculty-want-key-results-from-the-global-ai-faculty-survey-2025 Cited in: Ch 33.

[36] Digital Education Council. (2026). AI in higher education global survey 2026. https://www.digitaleducationcouncil.com/resource-library-items/ai-in-higher-education-global-survey-2026 Cited in: Ch 18, Ch 22, Ch 33.

[37] Digital Education Council. (2026). AI skills opportunity map. https://www.digitaleducationcouncil.com/resource-library-items/ai-skills-opportunity-map Cited in: Ch 30.

[38] Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. Science Advances, 10(28), eadn5290. https://doi.org/10.1126/sciadv.adn5290 Cited in: Ch 27.

[39] EdTech Innovation Hub. (2026). DfE publishes KCSIE 2026 with mobile phone-free default and AI safeguarding rules. https://www.edtechinnovationhub.com/news/dfe-publishes-kcsie-2026-with-mobile-phone-free-default-and-ai-safeguarding-rules Cited in: Ch 43.

[40] Education Endowment Foundation. (2024). ChatGPT in lesson preparation: A Teacher Choices trial. https://educationendowmentfoundation.org.uk/projects-and-evaluation/projects/choices-in-edtech-using-generative-ai-chatgpt-for-ks3-science-lesson-preparation-2024-teacher-choices-trial Cited in: Ch 34.

[41] Education International. (2025, December 5). Using artificial intelligence: Takeaways from TALIS 2024. https://www.ei-ie.org/en/item/31815:using-artificial-intelligence-takeaways-from-talis-2024 Cited in: Ch 33, Ch 48.

[42] eSchool News. (2026, March 25). Student use of AI for homework rises. https://www.eschoolnews.com/digital-learning/2026/03/25/student-use-of-ai-for-homework-rises/ Cited in: Ch 38.

[43] European Commission & OECD. (2026). Empowering learners for the age of AI: An AI literacy framework for primary and secondary education. https://ailiteracyframework.org Cited in: Ch 19.

[44] European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj Cited in: Ch 43.

[45] Fabo, C. (2026). Who educates the LLMs? The uncertain role of the nation-state. In UNESCO, The algorithm in the room (pp. 121–124). UNESCO. https://doi.org/10.54675/PMRO5023 Cited in: Ch 45.

[46] Future Science. (2026). Frequently asked questions. https://future-science.org/faq Cited in: Ch 54.

[47] Future Science. (2026). Future Science Archives: Information for authors. https://future-science.org/future-science-archives Cited in: Ch 54.

[48] Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. Societies, 15(1), 6. https://doi.org/10.3390/soc15010006 Cited in: Ch 10, Ch 15.

[49] Gibson Dunn. (2026). EU AI Act Omnibus agreement: Postponed high-risk deadlines and other key changes. https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/ Cited in: Ch 43.

[50] Goldin, C., & Katz, L. F. (2008). The race between education and technology. Harvard University Press. Cited in: Ch 31.

[51] GovTech. (2026). NYC schools prohibit AI for grading, discipline, IEPs. https://www.govtech.com/education/k-12/nyc-schools-prohibit-ai-for-grading-discipline-ieps Cited in: Ch 41.

[52] GovTech. (2026). Stanford researchers find thin evidence behind AI classroom tools. https://www.govtech.com/education/k-12/stanford-researchers-find-thin-evidence-behind-ai-classroom-tools Cited in: Ch 34, Ch 44.

[53] GSMA. (2025). The state of mobile internet connectivity 2025. https://www.gsma.com/somic/ Cited in: Ch 48.

[54] Guillet, B. D., & Penfold, P. (2013). Conducting immersive research in Second Life: A hotel co-branding case study. International Journal of Hospitality & Tourism Administration, 14(1). https://doi.org/10.1080/15256480.2013.753803 Cited in: Ch 52.

[55] Haarsbjerg, M. (2026). Between wood and algorithms: Crafting AI literacy in TVET. In UNESCO, The algorithm in the room (pp. 73–77). UNESCO. https://doi.org/10.54675/PMRO5023 Cited in: Ch 31.

[56] Inside Higher Ed. (2026, February 26). Agentic AI can complete whole courses for students. Now what? https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2026/02/26/agentic-ai-can-complete-whole-courses-now Cited in: Ch 24.

[57] Inside Higher Ed. (2026, March 23). Canvas unrolls AI teaching agent. https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2026/03/23/canvas-unrolls-ai-teaching-agent Cited in: Ch 24.

[58] International Labour Organization. (2026). Workers' exposure to AI: What indicators tell us, and what they don't (Research brief). ILO. https://www.ilo.org/sites/default/files/2026-03/Research%20Brief_Workers%20exposure%20to%20AI.pdf Cited in: Ch 29.

[59] International Rescue Committee. (2026). aprendIA in Nigeria. https://www.rescue.org/aprendIA/Nigeria Cited in: Ch 17, Ch 50.

[60] Kendeou, P., Greene, J. A., Nixon, N., et al. (2026). Ways generative AI can support and threaten learning in K–20 US education (Expert Consensus Report). University of Minnesota College of Education and Human Development, with the UNC-Chapel Hill and UC Irvine Schools of Education. https://cehd.umn.edu/AI-expert-consensus-report-2026 Cited in: Ch 7.

[61] Kennedy, M. M. (2016). How does professional development improve teaching? Review of Educational Research, 86(4), 945–980. https://doi.org/10.3102/0034654315626800 Cited in: Ch 35.

[62] Kestin, G., Miller, K., Klales, A., Milbourne, T., & Ponti, G. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. Scientific Reports, 15. https://doi.org/10.1038/s41598-025-97652-6 Cited in: Ch 2, Ch 3, Ch 10.

[63] Kokotajlo, D., Alexander, S., Larsen, T., Lifland, E., & Dean, R. (2025). AI 2027. https://ai-2027.com/ Cited in: Ch 3.

[64] Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task [Preprint]. arXiv. https://arxiv.org/abs/2506.08872 Cited in: Ch 10, Ch 15.

[65] Kraft, M. A., Blazar, D., & Hogan, D. (2018). The effect of teacher coaching on instruction and achievement: A meta-analysis of the causal evidence. Review of Educational Research, 88(4), 547–588. https://doi.org/10.3102/0034654318759268 Cited in: Ch 35.

[66] Larivière, V., Haustein, S., & Mongeon, P. (2015). The oligopoly of academic publishers in the digital era. PLOS ONE, 10(6), e0127502. https://doi.org/10.1371/journal.pone.0127502 Cited in: Ch 53.

[67] Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems. ACM. https://doi.org/10.1145/3706598.3713778 Cited in: Ch 10.

[68] Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. Patterns, 4(7), 100779. https://doi.org/10.1016/j.patter.2023.100779 Cited in: Ch 25.

[69] Liang, W., Zhang, Y., Wu, Z., Lepp, H., Ji, W., et al. (2025). Quantifying large language model usage in scientific papers. Nature Human Behaviour, 9(12), 2599–2609. https://doi.org/10.1038/s41562-025-02273-8 Cited in: Ch 53.

[70] Lodge, J. M., Howard, S., Bearman, M., Dawson, P., & Associates. (2023). Assessment reform for the age of artificial intelligence. Tertiary Education Quality and Standards Agency. https://www.teqsa.gov.au/sites/default/files/2023-09/assessment-reform-age-artificial-intelligence-discussion-paper.pdf Cited in: Ch 26.

[71] Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (pp. 1–16). ACM. https://doi.org/10.1145/3313831.3376727 Cited in: Ch 18.

[72] MacAskill, W., & Caviola, L. (2026, July 19). Could AI be conscious? The Guardian. https://www.theguardian.com/technology/2026/jul/19/could-ai-be-conscious Cited in: Ch 16.

[73] Ministry of Education of the People's Republic of China. (2024). Notice of the Office of the Ministry of Education on strengthening artificial intelligence education in primary and secondary schools (Center for Security and Emerging Technology, Trans.). https://cset.georgetown.edu/publication/china-primary-secondary-ai-education-notice/ Cited in: Ch 21.

[74] Muellauer, C. (2026). When my students told me I 'sound like AI'. In UNESCO, The algorithm in the room (pp. 69–72). UNESCO. https://doi.org/10.54675/PMRO5023 Cited in: Ch 14.

[75] Naffi, N. (2026). Deepfakes and the crisis of knowing. In UNESCO, The algorithm in the room (pp. 31–36). UNESCO. https://doi.org/10.54675/PMRO5023 Cited in: Ch 15.

[76] Namvarpour, M., Brofsky, D., Medina, J., Akter, M., & Razi, A. (2026). Understanding teen overreliance on AI companion chatbots through self-reported Reddit narratives. In Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems. ACM. https://arxiv.org/abs/2507.15783 Cited in: Ch 14.

[77] Narayanan, A., & Kapoor, S. (2025). AI as normal technology. Knight First Amendment Institute. https://knightcolumbia.org/content/ai-as-normal-technology Cited in: Ch 3.

[78] National Law Review. (2025). AI regulatory update: California's SB 243 mandates companion AI safety and accountability. https://natlawreview.com/article/ai-regulatory-update-californias-sb-243-mandates-companion-ai-safety-and Cited in: Ch 16.

[79] National Law Review. (2026). EU Digital Omnibus on AI enters into force. https://natlawreview.com/article/eu-digital-omnibus-ai-enters-force Cited in: Ch 43.

[80] Nature. (2025, November 7). Preprint site arXiv is banning computer-science reviews: Here's why. Nature [News]. https://www.nature.com/articles/d41586-025-03664-7 Cited in: Ch 2, Ch 53.

[81] Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. Computers and Education: Artificial Intelligence, 2, 100041. https://doi.org/10.1016/j.caeai.2021.100041 Cited in: Ch 18.

[82] Novik, A. (2026). Outpaced: When the world changes faster than children grow. In UNESCO, The algorithm in the room (pp. 45–49). UNESCO. https://doi.org/10.54675/PMRO5023 Cited in: Ch 15.

[83] NPR. (2026, January 27). In China, AI is no longer optional for some kids. It's part of the curriculum. https://www.npr.org/2026/01/27/nx-s1-5683821/china-ai-schools-curriculum Cited in: Ch 21.

[84] NPR. (2026, June 5). Poll: Teachers worry AI is impacting students' critical thinking. https://www.npr.org/2026/06/05/nx-s1-5779757/school-ai-education-students-teachers-poll-critical-thinking Cited in: Ch 24, Ch 33, Ch 34.

[85] NPR. (2026, September 28). Schools are experimenting with AI with little evidence or policy to guide them. https://news.wfsu.org/npr-news/2026-09-28/schools-are-experimenting-with-ai-with-little-evidence-or-policy-to-guide-them Cited in: Ch 44.

[86] OECD. (2020). Back to the future(s) of education: Four OECD scenarios for schooling. OECD Publishing. https://doi.org/10.1787/178ef527-en Cited in: Ch 1, Ch 3.

[87] OECD. (2025). Results from TALIS 2024: The state of teaching. OECD Publishing. https://www.oecd.org/en/publications/results-from-talis-2024_90df6235-en.html Cited in: Ch 33.

[88] OECD. (2026). Ecosystems approach to curriculum change (OECD Education Policy Perspectives). OECD Publishing. https://www.oecd.org/en/publications/ecosystems-approach-to-curriculum-change_df1310fd-en.html Cited in: Ch 40.

[89] OECD. (2026, April 9). International aid fell sharply in 2025, says OECD [Press release]. https://www.oecd.org/en/about/news/press-releases/2026/04/international-aid-fell-sharply-in-2025-says-oecd.html Cited in: Ch 2, Ch 49.

[90] OECD. (2026). OECD Digital Education Outlook 2026: Exploring effective uses of generative AI in education. OECD Publishing. https://www.oecd.org/en/publications/oecd-digital-education-outlook-2026_062a7394-en.html Cited in: Ch 41.

[91] OECD. (2026). PISA 2029 media and artificial intelligence literacy (MAIL) assessment framework: First draft. https://www.oecd.org/en/about/projects/pisa-2029-media-and-artificial-intelligence-literacy.html Cited in: Ch 19.

[92] Oltman, S. (2026, June 11). AI absolutism is breaking our brains. The Guardian. https://www.theguardian.com/technology/2026/jun/11/ai-absolutism-apocalyptic-future Cited in: Ch 16.

[93] O'Sullivan, J., & Kelly, M. (2026). Environmental contradictions of large language models in higher education. In UNESCO, The algorithm in the room (pp. 125–130). UNESCO. https://doi.org/10.54675/PMRO5023 Cited in: Ch 45.

[94] Penfold, P. (2009). Learning through the world of Second Life: A hospitality and tourism experience. Journal of Teaching in Travel & Tourism, 8(2–3), 139–160. https://doi.org/10.1080/15313220802634224 Cited in: Ch 52.

[95] Penfold, P. (2026). A new blueprint for learning: Navigating the K-12 AI shift. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/new-blueprint-learning-navigating-k-12-ai-shift-paul-penfold-zt8wc Cited in: Ch 37, Ch 39.

[96] Penfold, P. (2026). AI cheating, leaked papers, and marking errors: How exam protests went global. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/ai-cheating-leaked-papers-marking-errors-how-exam-protests-penfold-m9imc Cited in: Ch 23, Ch 26, Ch 27.

[97] Penfold, P. (2026). AI could close the gap or blow it wide open. Schools will help decide which. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/ai-could-close-gap-blow-wide-open-schools-help-decide-paul-penfold-396yc Cited in: Ch 47, Ch 48, Ch 51.

[98] Penfold, P. (2026). AI in 2035: Four futures for life and learning. Educational Futures, Part 5. https://www.linkedin.com/pulse/ai-2035-four-futures-life-learning-paul-penfold-sdflc Cited in: Ch 1, Ch 4, Ch 6.

[99] Penfold, P. (2026). AI is hiring more people, not fewer. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/ai-hiring-more-people-fewer-paul-penfold-w3joc Cited in: Ch 28, Ch 29.

[100] Penfold, P. (2026). AI is making students perform better and learn less. AI in Education – Deliberately Human newsletter, June 2026. https://paulpenfold.com/blog/ai-making-students-perform-better-learn-less-paul-penfold-jlt4c Cited in: Ch 12.

[101] Penfold, P. (2026). AI just became a safeguarding duty, not a tech decision. Most schools haven't noticed yet. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/ai-just-became-safeguarding-duty-tech-decision-most-schools-penfold-sn0sc Cited in: Ch 42, Ch 43, Ch 46.

[102] Penfold, P. (2026). AI will not decide the future of education – we will. Educational Futures, Part 1. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/ai-decide-future-education-we-paul-penfold-lo81f Cited in: Ch 1, Ch 3, Ch 6.

[103] Penfold, P. (2026). America's biggest school system just hit pause on AI. Here's what we can learn from it. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/americas-biggest-school-system-just-hit-pause-ai-heres-paul-penfold-if6pc Cited in: Ch 18, Ch 21, Ch 22.

[104] Penfold, P. (2026). Beyond automation: Reimagining the teacher's superpower. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/beyond-automation-reimagining-teachers-superpower-paul-penfold-mgdec Cited in: Ch 32, Ch 34.

[105] Penfold, P. (2026). Beyond cheating: Does the current assessment model still work? AI in Education – Deliberately Human. https://www.linkedin.com/pulse/beyond-cheating-does-current-assessment-model-still-work-paul-penfold-2b3qc Cited in: Ch 23, Ch 25.

[106] Penfold, P. (2026). Boston made AI fluency mandatory. One comment asked the harder question. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/boston-made-ai-fluency-mandatory-one-comment-asked-harder-penfold-dmgzc Cited in: Ch 18, Ch 21, Ch 22.

[107] Penfold, P. (2026). Childhood doesn't move at Silicon Valley speed. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/growing-up-machine-speed-paul-penfold-vkpwc Cited in: Ch 13, Ch 15, Ch 17.

[108] Penfold, P. (2026). China cuts 12,000 university degrees for the AI era. Here's the bet it's making. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/china-cuts-12000-university-degrees-ai-era-heres-bet-its-paul-penfold-n43me Cited in: Ch 28, Ch 31.

[109] Penfold, P. (2026). China just put a date on something most school systems are still debating in principle. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/china-just-put-date-something-most-school-systems-still-paul-penfold-4v6qc Cited in: Ch 18, Ch 20, Ch 21.

[110] Penfold, P. (2026). Empowering the educator: Why teacher agency is non-negotiable. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/empowering-educator-why-teacher-agency-non-negotiable-paul-penfold-3yrwc Cited in: Ch 32, Ch 36.

[111] Penfold, P. (2026). Europe just blinked on its AI rules for schools. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/europe-just-blinked-its-ai-rules-schools-paul-penfold-8rdvc Cited in: Ch 42, Ch 43, Ch 46.

[112] Penfold, P. (2026). From policy to practice: How to actually implement AI readiness. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/from-policy-practice-how-actually-implement-ai-paul-penfold-w7xlc Cited in: Ch 37, Ch 41.

[113] Penfold, P. (2026). How we're re-designing assessment for the AI age. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/how-were-re-designing-assessment-ai-age-paul-penfold-r2bmc Cited in: Ch 23, Ch 27.

[114] Penfold, P. (2026). If AI can produce the answer, what should schools assess? AI in Education – Deliberately Human. https://www.linkedin.com/pulse/ai-can-produce-answer-what-should-schools-assess-paul-penfold-hh2vf Cited in: Ch 1, Ch 23, Ch 27.

[115] Penfold, P. (2026). If education is a common good, who's actually paying? AI in Education – Deliberately Human. https://www.linkedin.com/pulse/education-common-good-whos-actually-paying-paul-penfold-rgkhc Cited in: Ch 42, Ch 45.

[116] Penfold, P. (2026). In a world of confident machines, critical thinking is the premium skill. AI in Education – Deliberately Human newsletter, October 2026. https://paulpenfold.com/blog/world-confident-machines-critical-thinking-premium-skill-paul-penfold-nbbjc Cited in: Ch 12.

[117] Penfold, P. (2026). Insights: AI in Education – Deliberately Human [Newsletter archive, 81 long-form articles]. https://paulpenfold.com/blog Cited in: Ch 54.

[118] Penfold, P. (2026). Learn first, then use the AI: What 30 experts say about generative AI and learning. AI in Education – Deliberately Human newsletter, October 2026. https://paulpenfold.com/blog/learn-first-use-ai-what-30-experts-say-generative-learning-penfold-ourdc Cited in: Ch 7, Ch 12.

[119] Penfold, P. (2026). Learning when schools, power and borders fail. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/learning-when-schools-power-borders-fail-paul-penfold-xwj8c Cited in: Ch 1, Ch 47, Ch 48, Ch 49, Ch 50.

[120] Penfold, P. (2026). Our students already use AI. The problem is we haven't taught them how. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/our-students-already-use-ai-problem-we-havent-taught-them-penfold-boooc Cited in: Ch 18, Ch 22.

[121] Penfold, P. (2026). Personalised to become what, exactly? AI in Education – Deliberately Human. https://www.linkedin.com/pulse/personalisation-myth-paul-penfold-tfuqc Cited in: Ch 13, Ch 14.

[122] Penfold, P. (2026). Plan for disruption, not a return to normal. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/plan-disruption-return-normal-paul-penfold-tc2mc Cited in: Ch 1, Ch 6.

[123] Penfold, P. (2026). Schools are buying AI on hope. Here's the evidence they're skipping. AI in Education – Deliberately Human newsletter, October 2026. https://paulpenfold.com/blog/schools-buying-ai-hope-heres-evidence-theyre-skipping-paul-penfold-d8twe Cited in: Ch 12, Ch 42, Ch 44, Ch 46.

[124] Penfold, P. (2026). Stop shopping for AI tools. Start redesigning the institution around them. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/stop-shopping-ai-tools-start-redesigning-institution-around-penfold-lisic Cited in: Ch 37, Ch 41.

[125] Penfold, P. (2026). The 39% problem: Why almost half of what we teach for work is about to change. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/39-problem-why-almost-half-what-we-teach-work-change-paul-penfold-uur7c Cited in: Ch 28, Ch 29, Ch 31.

[126] Penfold, P. (2026). The best AI skill you haven't trained for: Knowing when to say no. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/best-ai-skill-you-havent-trained-knowing-when-say-paul-penfold-yfm6c Cited in: Ch 18, Ch 20, Ch 32, Ch 36.

[127] Penfold, P. (2026). The chatbot your students use might have feelings. Now what? AI in Education – Deliberately Human. https://www.linkedin.com/pulse/chatbot-your-students-use-might-have-feelings-now-what-paul-penfold-7ixhc Cited in: Ch 13, Ch 16.

[128] Penfold, P. (2026). The classroom AI doom loop is already running. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/classroom-ai-doom-loop-already-running-paul-penfold-klnsc Cited in: Ch 23, Ch 24, Ch 27.

[129] Penfold, P. (2026). The Confidence Trap: Why flawless logic isn't always truth. AI in Education – Deliberately Human newsletter, June 2026. https://paulpenfold.com/blog/confidence-trap-why-flawless-logic-isnt-always-truth-paul-penfold-j22vc Cited in: Ch 12, Ch 13, Ch 15.

[130] Penfold, P. (2026). The Council of Europe just told 46 governments what AI literacy actually means. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/council-europe-just-told-46-governments-what-ai-literacy-paul-penfold-ntkoc Cited in: Ch 18, Ch 19, Ch 20.

[131] Penfold, P. (2026). The curriculum isn't the problem. The system around it is. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/curriculum-isnt-problem-system-around-paul-penfold-bnh1c Cited in: Ch 37, Ch 40.

[132] Penfold, P. (2026). The degree paradox: Why your credentials might be decoupling. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/degree-paradox-why-your-credentials-might-decoupling-paul-penfold-rfboc Cited in: Ch 28, Ch 31.

[133] Penfold, P. (2026). The empathy deficit: Is AI replacing the heart of the classroom? AI in Education – Deliberately Human. https://www.linkedin.com/pulse/empathy-deficit-ai-replacing-heart-classroom-paul-penfold-i947c Cited in: Ch 13, Ch 14.

[134] Penfold, P. (2026). The EU Council just spent its first-ever education session on AI. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/eu-council-just-spent-its-first-ever-education-session-paul-penfold-lfukc Cited in: Ch 32, Ch 36, Ch 42, Ch 43.

[135] Penfold, P. (2026). The governance challenge: Making schools safe for innovation. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/governance-challenge-making-schools-safe-innovation-paul-penfold-le3sc Cited in: Ch 46.

[136] Penfold, P. (2026). The great release: A better way to learn for our children and grandchildren. The Future of Learning by 2045, Part 3. Unpublished manuscript. Cited in: Ch 1, Ch 5, Ch 6.

[137] Penfold, P. (2026). The hidden infrastructure: Governance, data, and energy. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/hidden-infrastructure-governance-data-energy-paul-penfold-tdocc Cited in: Ch 42, Ch 45.

[138] Penfold, P. (2026). The last lecture: Will schools and universities still matter in 2045? The Future of Learning by 2045, Part 1. https://www.linkedin.com/pulse/last-lecture-schools-universities-still-matter-2045-paul-penfold-8xhrc Cited in: Ch 1, Ch 5, Ch 6.

[139] Penfold, P. (2026). The Lazy Brain: Why convenience is the enemy of intelligence. AI in Education – Deliberately Human newsletter, June 2026. https://paulpenfold.com/blog/lazy-brain-why-convenience-enemy-intelligence-paul-penfold-fhd3c Cited in: Ch 12, Ch 13, Ch 15.

[140] Penfold, P. (2026). The most important AI conversation your school hasn't had yet. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/most-important-ai-conversation-your-school-hasnt-had-yet-paul-penfold-xwmac Cited in: Ch 37, Ch 41.

[141] Penfold, P. (2026). The people building AI just told us to slow down. Education should listen. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/people-building-ai-just-told-us-slow-down-education-should-penfold-mvxtc Cited in: Ch 37, Ch 39, Ch 41.

[142] Penfold, P. (2026). The quiet crisis: AI is coming for the jobs where people learn. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/quiet-crisis-ai-coming-jobs-where-people-learn-paul-penfold-c4sic Cited in: Ch 28, Ch 29, Ch 30, Ch 31.

[143] Penfold, P. (2026). The robots are coming for routine work. Caring is booming. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/robots-coming-

routine-work-caring-booming-paul-penfold-ctjsc Cited in: Ch 28, Ch 31.

[144] Penfold, P. (2026). The skills gap is a message to schools. Are we listening? AI in Education – Deliberately Human. https://www.linkedin.com/pulse/skills-gap-message-schools-we-listening-paul-penfold-mf8bc Cited in: Ch 28, Ch 29, Ch 31.

[145] Penfold, P. (2026). The speed gap: Why institutions are racing against algorithms. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/speed-gap-why-institutions-racing-against-algorithms-paul-penfold-cg3zc Cited in: Ch 37, Ch 38.

[146] Penfold, P. (2026). The water bill behind your school's AI subscription. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/water-bill-behind-your-schools-ai-subscription-paul-penfold-4utxc Cited in: Ch 42, Ch 45, Ch 46.

[147] Penfold, P. (2026). Two schools, two futures: Who will afford a human teacher in 2045? The Future of Learning by 2045, Part 2. https://www.linkedin.com/pulse/two-schools-futures-who-afford-human-teacher-2045-paul-penfold-qtcwc Cited in: Ch 1, Ch 5, Ch 6.

[148] Penfold, P. (2026). Unbundling school: A global movement. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/unbundling-school-global-movement-paul-penfold-lxdjc Cited in: Ch 37, Ch 41.

[149] Penfold, P. (2026). What AI still can't do with its hands. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/what-ai-still-cant-do-its-hands-paul-penfold-wlhtc Cited in: Ch 28, Ch 31.

[150] Penfold, P. (2026). What does "AI literacy" actually mean? The Digital Education Council has an answer worth reading. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/what-does-ai-literacy-actually-mean-digital-education-paul-penfold-ddlac Cited in: Ch 18, Ch 19.

[151] Penfold, P. (2026). What if we just… went back to books? AI in Education – Deliberately Human. https://www.linkedin.com/pulse/what-we-just-went-back-books-paul-penfold-iko4c Cited in: Ch 17.

[152] Penfold, P. (2026). What is school actually for? Rethinking the purpose, not just the methods. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/what-school-actually-rethinking-purpose-just-methods-paul-penfold-fpaec Cited in: Ch 37, Ch 39.

[153] Penfold, P. (2026). When the machine learns to care. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/when-machine-learns-care-paul-penfold-jug7c Cited in: Ch 13, Ch 14, Ch 17.

[154] Penfold, P. (2026). Which tasks can AI actually do, and which ones stay human? AI in Education – Deliberately Human. https://www.linkedin.com/pulse/which-tasks-can-ai-actually-do-ones-stay-human-paul-penfold-a25lc Cited in: Ch 28, Ch 30, Ch 31.

[155] Penfold, P. (2026). Who benefits when AI enters the classroom? Educational Futures, Part 2. https://www.linkedin.com/pulse/who-benefits-when-ai-enters-classroom-paul-penfold-7iddf Cited in: Ch 1, Ch 3.

[156] Penfold, P. (2026). Why "AI will save us" and "AI will destroy us" are the same sales pitch. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/why-ai-save-us-destroy-same-sales-pitch-paul-penfold-hbdtc Cited in: Ch 16.

[157] Penfold, P. (2026). Why PISA 2029's new media & AI literacy test deserves your attention. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/why-pisa-2029s-new-media-ai-literacy-test-deserves-your-paul-penfold-7ol2c Cited in: Ch 18, Ch 19, Ch 22.

[158] Penfold, P. (2026). You can't teach what you're avoiding: Closing the teacher AI gap. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/you-cant-teach-what-youre-avoiding-closing-teacher-ai-paul-penfold-nryfc Cited in: Ch 32, Ch 33, Ch 35, Ch 36.

[159] Penfold, P. (2026). Your AI cheat detector is failing your most vulnerable students. AI in Education – Deliberately Human. https://www.linkedin.com/pulse/your-ai-cheat-detector-failing-most-vulnerable-students-paul-penfold-4mm8c Cited in: Ch 23, Ch 25.

[160] Priestley, M., Biesta, G., & Robinson, S. (2015). Teacher agency: An ecological approach. Bloomsbury. Cited in: Ch 36.

[161] PwC. (2026). 2026 Global AI Jobs Barometer. https://www.pwc.com/gx/en/news-room/press-releases/2026/pwc-2026-ai-jobs-barometer.html Cited in: Ch 29.

[162] Rahman, A., & Freeman, A. (2025). Artificial intelligence and education systems in 2035: Fourteen trends and five scenarios for how the future might unfold. EdTech Hub and IDRC. https://doi.org/10.53832/edtechhub.1125 Cited in: Ch 1, Ch 3.

[163] Rahman, A., & Freeman, A. (2025). Education in conflict and crisis in 2035: Twelve trends and five scenarios for how the future might unfold. EdTech Hub and IDRC. https://doi.org/10.53832/edtechhub.1126 Cited in: Ch 1, Ch 3, Ch 49, Ch 50.

[164] Roediger, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. Psychological Science, 17(3), 249–255. https://doi.org/10.1111/j.1467-9280.2006.01693.x Cited in: Ch 9.

[165] Sanchez Salas, M. (2026). Trauma, education and AI in conflict contexts: Care and control. In UNESCO, The algorithm in the room (pp. 54–58). UNESCO. https://doi.org/10.54675/PMRO5023 Cited in: Ch 17.

[166] Sanoma Learning. (2026, September 24). European teachers are adopting AI rapidly, but want tools built for education. GlobeNewswire. https://www.globenewswire.com/news-release/2026/09/24/3367985/0/en/european-teachers-are-adopting-ai-rapidly-but-want-tools-built-for-education.html Cited in: Ch 33, Ch 44.

[167] Save the Children. (2025, June 5). About 300,000 children risk losing education as learning centres in Rohingya camps shut due to funding cuts. https://www.savethechildren.net/news/about-300000-children-risk-losing-education-learning-centres-rohingya-camps-shut-due-funding Cited in: Ch 49.

[168] Save the Children. (2026, January 20). Education in refugee camps on Thailand–Myanmar border reaches "breaking point". https://www.savethechildren.net/news/education-refugee-camps-thailand-myanmar-border-reaches-breaking-point-report Cited in: Ch 49.

[169] SIPRI. (2026, April 27). Global military spending rise continues as European and Asian expenditures surge [Press release]. https://www.sipri.org/media/press-release/2026/global-military-spending-rise-continues-european-and-asian-expenditures-surge Cited in: Ch 2, Ch 49.

[170] Slamecka, N. J., & Graf, P. (1978). The generation effect: Delineation of a phenomenon. Journal of Experimental Psychology: Human Learning and Memory, 4(6), 592–604. https://doi.org/10.1037/0278-7393.4.6.592 Cited in: Ch 9.

[171] Solomon, D., & Holmes, O. (2026, August 16). AI cheating, leaked papers and marking errors: How exam protests went global. The Guardian. Syndicated copy: https://www.aol.co.uk/articles/ai-cheating-leaked-papers-marking-070019000.html Cited in: Ch 26.

[172] South China Morning Post. (2026). China's universities cut 12,000 'obsolete' degrees amid race to embrace AI era. https://www.scmp.com/economy/china-economy/article/3356913/chinas-universities-cut-12000-obsolete-degrees-amid-race-embrace-ai-era Cited in: Ch 31.

[173] Spence, M. (1973). Job market signaling. Quarterly Journal of Economics, 87(3), 355–374. https://doi.org/10.2307/1882010 Cited in: Ch 31.

[174] Spring University Myanmar. (n.d.). https://www.springuniversitymm.com/ Listed in the source paper; not cited in a chapter.

[175] Stanford HAI. (2026). The 2026 AI Index report: Education. Stanford Institute for Human-Centered Artificial Intelligence. https://hai.stanford.edu/ai-index/2026-ai-index-report/education Cited in: Ch 18, Ch 22, Ch 33.

[176] StateScoop. (2026). NYC schools ban student use of generative AI in classrooms through middle school. https://statescoop.com/nyc-schools-ban-student-use-of-generative-ai-in-classrooms-through-middle-school/ Cited in: Ch 21.

[177] Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. Cognitive Science, 12(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4 Cited in: Ch 9.

[178] Times Higher Education. (2026, May 14). Two-fifths of universities expect deficits amid student downturn. https://www.timeshighereducation.com/node/742942 Cited in: Ch 2.

[179] Tooley, J. (2009). The beautiful tree: A personal journey into how the world's poorest people are educating themselves. Cato Institute. Cited in: Ch 41.

[180] UK Government Office for Science. (2026). AI scenarios 2030: Helping policymakers plan for the future of AI. https://www.gov.uk/government/publications/ai-scenarios-2030-helping-policymakers-plan-for-the-future-of-ai/ai-scenarios-2030-helping-policymakers-plan-for-the-future-of-ai Cited in: Ch 1, Ch 3.

[181] UNDP. (n.d.). Kolibri (Learning Equality). Digital X. https://digitalx.undp.org/catalogs/kolibri.html Listed in the source paper; not cited in a chapter.

[182] UNESCO. (2021). Reimagining our futures together: A new social contract for education. UNESCO. https://doi.org/10.54675/ASRB4722 Cited in: Ch 1, Ch 3.

[183] UNESCO. (2023). Global education monitoring report 2023: Technology in education: A tool on whose terms? UNESCO. https://doi.org/10.54676/UZQV8501 Cited in: Ch 45.

[184] UNESCO. (2023). Guidance for generative AI in education and research. UNESCO. https://doi.org/10.54675/EWZM9535 Cited in: Ch 12.

[185] UNESCO. (2024). AI competency framework for students. UNESCO. https://doi.org/10.54675/JKJB9835 Cited in: Ch 19.

[186] UNESCO. (2024). AI competency framework for teachers. UNESCO. https://doi.org/10.54675/ZJTE2084 Cited in: Ch 19, Ch 36.

[187] UNESCO. (2024). Global report on teachers: Addressing teacher shortages and transforming the profession. UNESCO. https://www.unesco.org/en/articles/global-report-teachers-what-you-need-know Cited in: Ch 2, Ch 34.

[188] UNESCO. (2026). Education ministers call for education to remain a common good in the age of AI. https://www.unesco.org/en/articles/education-ministers-call-education-remain-common-good-age-ai-unescos-digital-learning-week Cited in: Ch 6, Ch 43.

[189] UNESCO. (2026). The algorithm in the room: Seeing and confronting the implications of AI and the futures of education. UNESCO. https://doi.org/10.54675/PMRO5023 Cited in: Ch 13, Ch 24, Ch 42, Ch 48.

[190] UNESCO-UNEVOC. (2026). Integrating AI in TVET: A practical guide for institutions (C. Yang, W. Wu, & H. Tegelbeckers). UNESCO. https://connect.unevoc.unesco.org/home/UNESCO+Publications/lang=en/akt=detail/qs=6849 Cited in: Ch 41.

[191] UNHCR. (2026, October 2). 5.8 million refugee children still out of school as funding cuts threaten hard-won progress. https://www.unhcr.org/news/press-releases/unhcr-5-8-million-refugee-children-still-out-school-funding-cuts-threaten-hard Cited in: Ch 2, Ch 49.

[192] UNICEF. (2025). After two years of war: Gaza's education system on the brink of collapse. https://www.unicef.org/sop/stories/after-two-years-war-gazas-education-system-brink-collapse Cited in: Ch 49.

[193] UNICEF. (2025, January 24). Nearly a quarter of a billion children's schooling was disrupted by climate crises in 2024. https://www.unicef.org/press-releases/nearly-quarter-billion-childrens-schooling-was-disrupted-climate-crises-2024-unicef Cited in: Ch 2.

[194] UNICEF. (2026). UNESCO, UNICEF and ITU launch landmark Charter for Public Digital Learning Platforms. https://www.unicef.org/digitaleducation/Charter-Launch Cited in: Ch 43, Ch 46.

[195] United Nations. (2024). World population prospects 2024: Key messages. https://population.un.org/wpp/assets/Files/WPP2024_Key-Messages.pdf Cited in: Ch 2.

[196] United Nations University. (2026). Environmental cost of AI's energy use: Carbon, water and land footprints. UNU-INWEH. https://unu.edu/inweh/news/environmental-cost-of-AIs-Enrgy-Use-carbon-water-and-land-footprints Cited in: Ch 45.

[197] University of Ottawa. (2026). uOttawa study finds AI use for emotional support may signal youth distress. https://www.uottawa.ca/en/news-all/uottawa-study-finds-ai-use-emotional-support-may-signal-youth-distress Cited in: Ch 14.

[198] WBUR. (2026, March 26). With new program, Boston to ensure AI literacy in public high schools. https://www.wbur.org/news/2026/03/26/boston-public-schools-ai-literacy Cited in: Ch 21.

[199] Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P., & Waddington, L. (2023). Testing of detection tools for AI-generated text. International Journal for Educational Integrity, 19, 26. https://doi.org/10.1007/s40979-023-00146-z Cited in: Ch 25.

[200] Wineburg, S., & McGrew, S. (2019). Lateral reading and the nature of expertise: Reading less and learning more when evaluating digital information. Teachers College Record, 121(11), 1–40. https://doi.org/10.1177/016146811912101102 Cited in: Ch 22.

[201] World Bank. (2026). AI and the future of education: From policy to implementation at scale (Global EdTech Policy Academy, Seoul, 31 August–4 September 2026). https://www.worldbank.org/en/events/2026/08/31/ai-future-education-policy-to-implementation Cited in: Ch 43.

[202] World Bank. (2026). World development report 2026: The promise of artificial intelligence. https://www.worldbank.org/en/publication/wdr2026 Cited in: Ch 50.

[203] World Economic Forum. (2025). Future of jobs report 2025. https://www.weforum.org/publications/the-future-of-jobs-report-2025/ Cited in: Ch 29, Ch 31.

[204] World Economic Forum. (2026). Four futures for jobs in the new economy: AI and talent in 2030. https://www.weforum.org/publications/four-futures-for-jobs-in-the-new-economy-ai-and-talent-in-2030/ Cited in: Ch 1, Ch 3.

[205] World Economic Forum. (2026). Shaping the future of learning: Education readiness for the age of AI. https://www.weforum.org/publications/shaping-the-future-of-learning-education-readiness-for-the-age-of-ai/ Cited in: Ch 13, Ch 14, Ch 15, Ch 25, Ch 26, Ch 27, Ch 31, Ch 34, Ch 36, Ch 38, Ch 39, Ch 41, Ch 43, Ch 45, Ch 46.

[206] World Economic Forum. (2026). Unbundling school: How a new financing mechanism is building modular education for the future. https://www.weforum.org/stories/education-and-skills/unbundling-school-how-a-new-financing-mechanism-is-building-modular-education-for-the-future/ Cited in: Ch 41.

[207] World Economic Forum & PwC. (2026). Artificial intelligence and the future of entry-level work: A framework for safeguarding and reinventing early career pathways. https://www.weforum.org/publications/artificial-intelligence-and-the-future-of-entry-level-work-a-framework-for-safeguarding-and-reinventing-early-career-pathways/ Cited in: Ch 29, Ch 30.

[208] Yidan Prize Foundation. (2026). 2026 Yidan Prize Conference Series: Seoul edition. https://yidanprize.org/events-and-news/events/2026-yidan-prize-conference-seoul Cited in: Ch 39.

[209] UNESCO, International Labour Organization, UNICEF & Education International. (2026, 5 October). World Teachers' Day joint message on the 60th anniversary of the Recommendation concerning the Status of Teachers. As reported by Africanews, 6 October 2026: https://www.africanews.com/2026/10/06/world-needs-50-million-additional-teachers-by-2030-unesco-says/ Cited in: Ch 2, Ch 34.

About the author

Paul Penfold is an Education and Technology Consultant based in Chiang Mai, Thailand, with six decades of experience in education, international development and humanitarian work.

He was Head of School at British Concordance International School in Thailand from 2023 to 2026, and has led education and training projects for governments and agencies in Asia, Europe and the Caribbean, including work with the European Union, UNWTO, Luxembourg Development Cooperation and the Asian Development Bank. He holds a Master in Education from the University of Southern Queensland, and is a Chartered Fellow of the CIPD and a Fellow of the RSA.

He is now Regional Coordinator for Brave Generation Academy in Thailand and Myanmar, advises schools and organisations on AI and learning, and writes the AI in Education – Deliberately Human newsletter.

More at https://paulpenfold.com/

Contact

Questions, invitations to speak or work together, and comments on the book are welcome.

Email: info@paulpenfold.com

Or use the contact form at https://paulpenfold.com/

Corrections and updates (errata)

This is an evolving book. Evidence on AI in education changes quickly, so some passages will be updated and some will turn out to be wrong. Every correction that changes a fact, a figure or a conclusion will be listed here with its date and the edition it applies to.

No corrections have been recorded yet for the first edition, October 2026.

If you spot an error, please email info@paulpenfold.com with the chapter, the sentence and, if you can, a source. The online edition is always the most current version.