

Measured in verbs.

In the age of AI, *what* you learn is being commoditised. *How* you learn is everything. Why the next era of personalised learning depends on making the process visible — one verb at a time.



Juliette Denny

Learning scientist · Author of *Learning to be Human* · Founder, Growth Engineering & Zavmo

The argument in one page

For a century we have measured learning by its outputs — courses completed, scores achieved, boxes ticked. We have measured *what* a person was taught, and almost never *how* they learned it. In the age of AI, that blind spot becomes a strategic failure. This paper argues that the future of personalised learning depends on making the *process* of learning visible — and that the technology to do it, the xAPI standard, has been sitting in plain sight for a decade.

- 1 **AI commoditises the "what."** When any fact or first draft is a prompt away, transferring knowledge stops being the scarce, valuable thing. What remains scarce is the *capacity to learn* — and that lives in the *how*, not the *what*.
- 2 **Personalisation is the proven prize — and it now scales.** Bloom showed one-to-one tutoring could lift the average learner far above the class.¹ Later evidence is more sober, but reveals something remarkable: intelligent tutoring systems are already *almost as effective as human tutors*.² The ceiling was never that the tutor was human.
- 3 **You cannot personalise what you cannot see.** A great tutor adapts by reading how you think — your hesitation, your misconception, your moment of mastery. An AI can only do the same if it can observe those signals. The signal is behaviour, and behaviour is captured as a *verb*.
- 4 **xAPI turns learning into a stream of verbs.** The Experience API records learning as *actor — verb — object* statements: *attempted, struggled, retried, connected, mastered*.³ It captures how someone learns, anywhere, not just whether they finished a course. Bloom's own taxonomy already describes thinking as verbs.⁵
- 5 **The how is also the highest-leverage thing to teach.** Formative, process-focused feedback is among the largest effects in education,⁶ and teaching learners to manage their own learning adds roughly seven months' progress a year.⁷ Measuring the how is how you finally grow the human fluencies no exam can certify.

The chapters that follow trace the argument from the economics of AI to the science of tutoring, to the standard that makes the invisible visible — and to what it means for anyone trying to help people learn.

SECTION 01

The question we stopped asking

Ask an organisation how its people are learning and you will get numbers: completion rates, pass marks, hours logged, courses consumed. Every one of them answers a

different question — *what* was delivered — while quietly ignoring the one that actually predicts growth: *how* did this person learn?

This is not a small omission. Two learners can finish the same module with the same score and have had completely different experiences: one breezed through material they already knew; the other wrestled with it, got it wrong, tried again, and finally understood. The score records neither the boredom nor the breakthrough. It is a photograph of a destination that tells you nothing about the journey — and it is the journey that decides whether the learning lasts, transfers, or evaporates by Friday.

We built our whole measurement apparatus this way for an understandable reason: the *how* was invisible. You could not see inside a learner's process at scale, so we measured what we could count. But that constraint has now been lifted twice over — once by the standards that let us capture the process, and once by the AI that makes measuring the *what* almost pointless. The rest of this paper is about what follows when you stop measuring the destination and start reading the journey.

A test score is a photograph of the destination. It tells you nothing about how the person travelled — and the travelling is the part that lasts.

SECTION 02

The personalisation prize — and why it eluded us

The dream of learning that adapts to the individual is not new, and it is not vague. It has been measured — and it works.

In 1984, the educational psychologist Benjamin Bloom published one of the most cited findings in the field. Students taught one-to-one, with mastery learning, performed around **two standard deviations** better than students in a conventional classroom — the average tutored learner outperforming some 98% of the conventionally taught group.¹ Bloom called it the "2 Sigma Problem," and the problem was economic, not pedagogical: one-to-one tutoring is wonderful and completely unaffordable at scale. His challenge to the field was to find a way to deliver that personalised effect to everyone.

For forty years, that challenge mostly went unmet — and the honest modern evidence is more sober than the legend. A careful meta-analysis put the real effect of skilled human tutoring nearer **0.8 standard deviations** than two.² But buried in the same analysis was a genuinely startling result: **intelligent tutoring systems — software — were almost as effective as human tutors** (an effect of roughly 0.76 versus 0.79).²

2σ

Bloom's benchmark for 1:1 mastery tutoring vs the classroom¹

0.79

Effect size of skilled human tutoring, on sober re-estimation²

0.76

Effect size of intelligent tutoring systems — nearly matching humans²

Sit with that finding, because it relocates the whole problem. If a machine tutor can nearly match a human one, then the thing that made tutoring powerful was never that the tutor was a person. It was that a tutor *watches you think* — and adjusts. A good tutor notices the flicker of confusion, the confident-but-wrong answer, the concept that finally clicks, and personalises the very next move accordingly. Whether that responsiveness comes from a human or a well-built system matters far less than whether it can happen at all. Which raises the only question that matters for scaling personalisation: **how does the tutor — human or AI — actually see how you are learning?**

SECTION 03

Why AI is commoditising the "what"

There is a reason this question has become urgent now, rather than at any point in the previous four decades. Generative AI has quietly demolished the value of the thing our education systems are built to deliver: the transfer of content.

For all of history, knowing things was scarce and therefore valuable. The expert was the person with the information; teaching was the act of moving that information into another head; assessment checked whether it had arrived. Generative AI collapses that economy. When any fact, explanation, summary or first draft is a sentence away, the marginal value of simply *possessing* the content falls towards zero. The "what" — the knowledge itself — is becoming a commodity.

What does not become a commodity is the human capacity to learn, judge, adapt and apply. The World Economic Forum finds that **39% of the core skills** workers rely on are expected to change by 2030, and it ranks **curiosity and lifelong learning** among the fastest-rising capabilities employers want.⁹ In a companion paper, *Future Fluency*, I argue that the capabilities the AI economy rewards most are the deeply human ones — and that they are precisely the ones no qualification can certify. This paper is the other half of that argument: if the durable value is in *how* people learn and adapt, then how they learn is the thing we must finally be able to see, develop and evidence.

The shift in one line. When knowledge was scarce, it made sense to measure whether it had been transferred. Now that knowledge is abundant, the only measurement worth making is of the learner's process — because the process is what AI cannot hand them, and what the modern economy actually pays for.


SECTION 04

The how is a behaviour — and a behaviour is a verb

If we want to see how someone learns, we need a way to record the acts of learning themselves — not the final score, but the attempts, the retries, the connections, the

moments of struggle and mastery. That record already exists as an open standard, and it has for over a decade.

The **Experience API** — xAPI, once nicknamed "Tin Can" — was released in 2013 by the Advanced Distributed Learning initiative, the same body that gave the world the SCORM e-learning standard.^{3,4} Where SCORM could only really record that a course had been completed and scored inside a learning management system, xAPI records *experiences*, anywhere, in a simple and profound grammar: **actor — verb — object**. "I *did* this." Each statement is written to a Learning Record Store, and it need not begin or end in a course at all.³



The verb is the heart of it. A completion is a single, silent full stop; a stream of verbs is a narrative. And notice that the richest verbs are not about *what* was covered but about *how* the learner behaved:



The raspberry verbs describe a process you can coach. The two pale verbs are the only ones a traditional system ever recorded.

This is not a coincidence of software design; it echoes how learning scientists already describe thinking. When Bloom's taxonomy was revised in 2001, its authors deliberately recast the levels of cognition as **action verbs** — remember, understand, apply, analyse, evaluate, create.⁵ To learn is to *do* something cognitive, and the verb names which cognitive act it was. xAPI simply gives us a way to record those acts as they happen. Put the two together and every meaningful moment of learning becomes a data point about *how* a person is thinking — not merely whether they reached the end.

SECTION 05

From verbs to hyperpersonalisation

Here the threads tie together. A verb stream is not just a richer record for humans to admire after the fact — it is the sensory input an AI needs to personalise in the moment, the way a great tutor does.

Recall the finding from Section 2: a machine tutor can nearly match a human one.² The reason a human tutor is effective is that they are continuously reading the learner's process and adjusting. An AI can only do the same if it receives the same signal — and a score at the end of a module is not that signal. A live stream of verbs is. When the system can see that a learner *attempted*, *hesitated*, *retried*, and *self-corrected*, it can infer the misconception and reshape the next question; when it sees *connected* and

applied, it can raise the challenge. **Hyperpersonalisation is not a feature you add to content; it is what becomes possible once the machine can observe how the learner is actually learning.**

The science of what to do with that signal is well established. Feedback that is *formative* — responsive to the learner's process rather than delivered as a final verdict — produces some of the largest effects in all of education research, with gains far above most interventions.⁶ And the most valuable thing the process data lets us build is the learner's own ability to manage their learning: teaching **metacognition and self-regulation** adds roughly **seven months of additional progress** a year, at low cost, on strong evidence.⁷ Self-regulated learning is itself a cycle of *planning, doing, and reflecting*⁸ — a loop of verbs. You cannot help a learner get better at that loop unless you can see them running it.

+7 mo

Progress a year from teaching learners to manage their own learning⁷

0.4–0.7

Effect-size range of formative, process-focused feedback⁶

1

The unit of the how: a single verb, captured as it happens³

A note on measuring responsibly. Reading how people learn is a form of trust, not surveillance. Behavioural learning data should be gathered to serve the learner — to coach, adapt and encourage — never to rank or police them, and always with transparency about what is captured and why. The point of seeing the struggle is to support it, not to punish it.

SECTION 06

What the verbs reveal that scores hide

To see why the how matters more than the what, put two learners side by side. They finish the same module. They earn the same final score. On every dashboard in use today, they are identical. Their verbs tell opposite stories.

Learner A

FINAL SCORE: 80%

attempted → *answered* → *completed*. Straight line, first time, no hesitation. The material was at or below what they already knew. Nothing was stretched; little new was built. The score is real, and the learning is shallow.

Learner B

FINAL SCORE: 80%

attempted → *struggled* → *retried* → *self-corrected* → *connected* → *mastered*. A hard, effortful path with a genuine breakthrough. The struggle is the sign of real, durable learning being formed — exactly the "desirable difficulty" the science prizes.

Same score; entirely different learning. Learner A should be moved on and challenged harder. Learner B has just done the most valuable thing a learner can do and should be recognised for it — yet a completion-based system would, if anything, flag B as the weaker student for taking longer. The verbs invert the judgement. They reveal effort, resilience, misconception and mastery — the very qualities that

predict whether someone will still be able to use this in a year, and the very qualities the AI economy most rewards.

This is the quiet revolution in measuring the how: it lets us finally value the things we always said mattered but could never see. Persistence. Curiosity. The willingness to be wrong on the way to being right. Not as soft aspirations, but as observable, coachable behaviour.

SECTION 07

Closing the loop with Future Fluency

This is where measuring the how meets the argument of the companion paper — and completes it.

In *Future Fluency* I set out the ten fluencies the AI economy is converging on, and drew a fault line through them: half can be taught and certified by qualifications — AI, data, governance — and half cannot. The human fluencies — growth mindset, systems thinking, foresight, the ability to take people through change — are learned by doing, in role, over time. A written exam cannot evidence them, because they are not knowledge you recall; they are behaviours you enact.

But behaviours are exactly what a verb stream captures. You cannot examine someone's resilience — yet you can observe that they *struggled, retried and persisted*. You cannot test for a growth mindset on paper — yet you can see a learner seek *harder challenges* and *treat a wrong answer as information*. Measuring the how is the missing instrument for the human half of Future Fluency: it is how we finally **develop, recognise and evidence the capabilities that no qualification can certify**. The two papers are one argument. Future Fluency says *these human capabilities are what matter most now*. This paper says *here is how we can finally see and grow them* — by watching how people learn, not just what they were taught.

You cannot certify a growth mindset on paper. But you can watch a learner reach for the harder problem — and that verb tells you everything.

SECTION 08

What this means in practice

Shifting the focus from what to how is not only a philosophy; it changes what good looks

like for anyone who designs, buys or delivers learning.

For L&D and educators — instrument the process, and prize the struggle

Stop treating completion and final scores as the measure of learning; they are the least informative data you hold. Design learning that emits meaningful verbs, and read them. Reward the learner who struggled and broke through, not only the one who finished fastest. The goal of measurement shifts from proving what was covered to understanding — and improving — how people actually learn.

For anyone building AI into learning — feed the model behaviour, not just answers

An AI tutor is only as personalised as the signal it receives. Right/wrong at the end of a module is a starvation diet. A live stream of process verbs is the input that lets an AI adapt the way an effective human tutor does — which is the only reason the technology can approach the tutoring effect at all.² Build on open standards like xAPI so the learner's process, not just their score, is portable and legible.

For individuals — your how is your edge

When AI can supply any answer, your durable advantage is how well you learn, adapt and apply — your metacognition, your self-regulation, your appetite to struggle productively.⁷ These are learnable, and they compound. The learner who understands *how* they learn will out-adapt the one who only knows more.

The instruments exist. The science is settled enough to act on. What has been missing is the decision to point our attention at the right thing — to stop counting what we deliver, and start understanding how people grow. In the age of AI, that is not a nicety. It is the whole game.

A CLOSING NOTE

Watch how, not what

I have always believed that people are capable of far more than the systems around them assume. For most of my career, the tragedy was that we could not see them growing — so we measured the shadow instead of the thing, and called a completed course a learned skill.

We no longer have that excuse. We can watch how people learn, in their own verbs, and use what we see to help them learn better. In an age where machines can hand us any answer, the most human thing we can do is get better at the one thing they cannot do for us: learning itself. That starts with paying attention to how.

Where this thinking is being built

These ideas are built into **Zavmo**, an AI learning companion I founded that pays attention to *how* people learn — coaching in the flow of real work, adapting to the learner's process rather than just their score. If you're rethinking how your organisation measures and grows its people, that's the conversation I'd love to have.

Read more, and the companion paper *Future Fluency*, at juliettedenny.com



About the author

Juliette Denny is a learning scientist and the author of *Learning to be Human*. She is the founder of Growth Engineering, an award-winning learning-technology company, and of Zavmo, an AI-native learning companion built on the science of how people actually learn. Her work sits at the intersection of pedagogy, behavioural science and technology, with a single throughline: designing learning that treats people as capable of growth.

REFERENCES & SOURCES

- 1 Bloom, B. S., *The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring*, Educational Researcher, 13(6), 1984. One-to-one mastery tutoring raised the average learner ~2 standard deviations (~98th percentile) above conventional instruction.
- 2 VanLehn, K., *The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems*, Educational Psychologist, 46(4), 2011. Human tutoring effect size ≈ 0.79 ; intelligent tutoring systems ≈ 0.76 — "nearly as effective as human tutoring."
- 3 Advanced Distributed Learning (ADL), *Experience API (xAPI)* — actor-verb-object statements written to a Learning Record Store; records learning experiences beyond the LMS. Overview at xapi.com; specification now IEEE 9274.1.1 (2023).
- 4 Advanced Distributed Learning (ADL), SCORM and the "Project Tin Can" origin of xAPI (v1.0, 2013). SCORM, xAPI's predecessor, recorded course completion and score within an LMS.
- 5 Anderson, L. W. & Krathwohl, D. R. (eds.), *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy*, 2001. The revised taxonomy recast the cognitive processes as action verbs: remember, understand, apply, analyse, evaluate, create.
- 6 Black, P. & Wiliam, D., *Inside the Black Box: Raising Standards Through Classroom Assessment*, Phi Delta Kappan, 1998. A review of ~580 studies found formative assessment among the largest effects in education (typical effect sizes ~ 0.4 – 0.7).
- 7 Education Endowment Foundation, *Teaching and Learning Toolkit — Metacognition and self-regulation*. Rated as adding roughly +7 months' additional progress a year, at low cost, on high-security evidence.
- 8 Zimmerman, B. J., cyclical model of *self-regulated learning* — a loop of forethought (planning), performance, and self-reflection, grounded in social-cognitive theory.
- 9 World Economic Forum, *Future of Jobs Report 2025* (2025). 39% of workers' core skills expected to change by 2030; "curiosity and lifelong learning" among the fastest-rising skills.
- 10 Denny, J., *Future Fluency: the UK skills economy in the age of AI* (2026) — companion white paper on the ten fluencies and the human capabilities the AI economy rewards. juliettedenny.com/papers.

© 2026 Juliette Denny. This white paper may be shared and cited with attribution. Effect sizes and findings are drawn from the peer-reviewed and institutional sources above; figures were current at the time of writing (July 2026). juliettedenny.com