

Learning Before Answering

An AI tutor could reach every Grade 9 student in Mauritius in 2027. We have been building around one question: how do you help without doing the learning for the learner?

ENTRY № 08 · 22 September 2026 · 15 min read · From the series *The Next Era of Education*

In the [previous essay](#) I argued that AI is arriving inside the school we already have, that this is useful, and that its first danger is that AI can make a student's work better before it makes the student better. I ended with a test. The measure of success is not how impressive the assisted answer becomes. It is how much more capable the learner becomes when the assistance is gone.

This essay is about applying that test to something real. In 2027, myt Learn could be rolled out to every Grade 9 student in Mauritius. myt Learn is a Mauritius Telecom product, and I am part of the team building it, so read what follows as an account from inside, not a review. I am writing it because the only honest way to prepare to put an AI tutor in front of a whole grade, having just published what can go wrong with AI tutors, is to say in public what this one is built to do, how it decides what to do next, and where the open questions still are.

The bet

Grade 9 is the year of the National Certificate of Education in Mauritius. It is the year in which the compromise of mass schooling, described in the first essay, is felt most sharply. Thirty learners, one adult, a syllabus that has to be covered by a fixed date, and an exam at the end that assumes everyone arrived at the same place at the same time. The learner who needed one more example does not get it. The learner who understood last week waits.

myt Learn does not change any of that. The syllabus, the textbooks, the classes and the exam are all exactly where they were. What it adds is a layer on top: a tutor that responds to each learner, tools built around teachers' work, and a view for parents. This is the third era in the [framework](#), help within the existing system, and I want to be clear that we chose it deliberately. Children sitting the NCE in 2027 cannot wait for a redesigned school. They can be helped through the school they are in.

The bet, then, is narrow and testable. Can a layer on top of the current system reduce the mismatch between what a fourteen-year-old needs and what the timetable can give, without producing the thing the previous essay warned about?

What the pilot changed

We did not start with a national ambition. myt Learn has been in pilot in secondary schools for more than a year, and is now in its second phase. The product Grade 9 students would use is the fourth version, and each version was shaped by what the

pilot showed us the previous one got wrong. When we say we want to deploy AI responsibly, this is what we mean in practice: a small deployment first, a long look at how it is actually used, changes, and then a wider deployment of one grade only, while pilots continue at every other level of secondary school.

The biggest thing the pilot changed was who the product was for. The early versions put most of our attention on the student. The pilot made clear that we had not given enough attention to the teacher's role in making the system useful. A tool placed between a teacher and their pupils, without tools for the teacher, does not get carried into classrooms, however good it is for the pupils. The response was not to explain the technology more persuasively. It was to change the product. Teachers needed tools designed around their work, not merely access to a tool designed for their pupils. The second phase built them.

I am not going to use the pilot numbers here as evidence that myt Learn improves learning. That was not what the pilot was designed to establish. Its purpose was to expose weaknesses in the product and change the next version. A wider rollout would pose a different question, and that deserves its own piece once there is real evidence to report.

The product that has emerged can be described through five commitments. Stay grounded in the curriculum. Respond to the learner's developing understanding. Preserve the work of learning. Strengthen human judgement. Protect the learner. The rest of this essay takes them in turn, and then says what has been tested and what none of it guarantees.

Stay grounded in the curriculum

The line the team uses internally is that a search engine answers questions and myt Learn teaches. Before that can mean anything, there is a prior question: teaches what?

The answer is the syllabus the student is actually on. myt Learn is grounded in the official textbooks for every subject at secondary level, digitised in full and used, with permission, to ground the tutor. The books remain the property of the institutions that publish them. A general model can give a fluent, broadly correct explanation of a Grade 9 chemistry idea that is still pitched at the wrong level, uses different terms from the lesson, or wanders into material the class will not meet for two years. Fluency and alignment are separate problems. A tutor grounded in the textbook answers from the chapter the class is on, using the terminology, sequence and level of the curriculum the student is studying.

The textbook is not only behind the tutor. It is available beside it. Each unit carries a references panel with the official textbook in digital form, the formulas and rules, key terms, worked examples, a summary, a map of the unit, and a section on common mistakes and how to avoid them. The learner has something to consult beyond the conversation. And learning is organised around units and sub-units of the curriculum rather than open prompting, which is a large part of how a chat interface becomes a course.

Respond to the learner, not just the question


Knowing what the curriculum contains is only half the problem. The other half is knowing where this learner is within it.

That is harder than it sounds, because answers are noisy evidence. A correct answer does not always mean understanding. It may be a guess. A wrong answer does not always mean the concept is missing. It may be a slip. And a learner who gets there after two hints and a worked example has shown something different from a learner who solves a new problem without help.

This is why myt Learn includes a method called Bayesian Knowledge Tracing. It maintains an evolving estimate of how securely a learner understands a particular skill, and updates that estimate as answers arrive. It allows for the possibility of a lucky answer or a careless mistake rather than treating every response as a verdict. The estimate is updated from the attempts that tell us most: the learner's first answer to a question, and their answer to a fresh question after help. An answer eventually reached through hints is useful for teaching, but it is not treated as equivalent evidence of independent understanding. That estimate is used to decide what happens next: whether to offer more practice, change the difficulty, or move forward to the next step. It is not a reading of the child's mind. It is a working estimate that can be wrong and must remain open to correction.

The next step depends on what the learner shows

How a learner's response changes what happens next



↪ Later review feeds new evidence back into the estimate, and the loop continues

□ Steps where the learner must act without help

The system estimates understanding. The learner still has to demonstrate it.

The aim is not to assign a label to the child. It is to make a better decision about the next learning step. And the estimate alone is not enough to complete a step. Eventually reaching the answer is not the same as demonstrating that you can now solve the problem, so the system distinguishes progress made with help from evidence of

understanding without it. That distinction is the practical answer to the previous essay. A tutor that could not tell the two apart would be manufacturing the illusion of learning at scale.

Preserve the work of learning


Here the previous essay’s argument becomes a design. But the design is more nuanced than “never give the answer”, and I want to be precise about it, because the nuance matters to parents and teachers.

A learner meeting an unfamiliar concept may need an explanation or a worked example before being asked to use it. Withholding it would be perverse. So in learning mode the tutor introduces the concept, step by step, before asking anything. A learner practising that concept needs space to attempt it. And a learner who has been helped through a problem needs a fresh opportunity to show what they can now do. The purpose is not to withhold explanations. It is to avoid mistaking an explanation the learner has seen for a capability the learner has acquired.

Nor is there a single teaching sequence for every subject. Mathematics may begin with a worked example and move through a guided attempt to independent practice. A conceptual topic may start with a check of what the learner already thinks, then an explanation, then the learner explaining it back. Writing needs a plan, a draft, feedback and a revision. The architecture is common. The pedagogy changes with what is being learned.

How help leads back to the learner

One practice question. A correct answer at any attempt exits this help sequence



If the fresh question is missed, the learner gets one more explanation and one more fresh question. After that the step stays incomplete, and the teacher can see it.

For each practice question, that is the chain. The learner tries. After a miss, a hint, and another try. After a second miss, a second hint, and a third try. After a third, a worked solution. Then, before the step can count, a new question with no hint at all. A correct answer at any attempt exits the sequence. Missing the fresh question earns one more explanation and one more fresh question, and after that the step stays incomplete and is visible to the teacher as such. The rule is not “make help difficult to obtain”. It is “make sure help leads back to the learner doing the thinking”.

The rule is not to make help hard to get. It is to make sure help leads back to the learner doing the thinking.

The same principle governs the open chat, where a learner can ask about anything on the syllabus. There the tutor withholds the final answer and explains the concept instead, which is the Socratic behaviour the previous essay described. Open in topic, closed in behaviour. That is the one place where the [Turkey study](#) is unavoidable: the unguarded interface left students about 17% behind those with no AI at all on the unaided exam, while the guarded tutor's students were indistinguishable from them. We did not want the unguarded version living one tab away from the guarded one.

Learning also has to survive the session. A successful session should not be the end of the tutor's interest in that concept. The next question is whether the learner can still use it later. So when a learner completes a sub-unit, the tutor returns to it the next day with three questions and no hints. If the learner remembers, the next return is further away. If not, it comes back tomorrow. Over weeks, the interval widens for what is secure and narrows for what is slipping. Revision works in the same direction. Rather than asking every learner to revise everything equally, the platform ranks a learner's units from weakest to strongest on their own exercise and quiz results, and can build a revision session aimed at the weakest, with a revision tutor, quizzes and flashcards to work through it.

I want to be careful not to overclaim. Scaffolding, fresh questions and later recall are the conditions for a tutor that builds capability rather than the appearance of it. They are not proof that it does. That proof can only come from the learners themselves, over time.

Strengthen human judgement

The pilot's lesson was that teachers needed tools designed around their work. Here is what that means, described through the decisions the tools support.


Marking. A teacher uploads submissions, including photographs of handwritten work, together with their own marking scheme and, if they wish, instructions on how strictly to apply it and how to allocate partial credit. The AI applies that framework to each script and returns marks and comments, question by question. It does not generate its own criteria. The teacher does not hand over the definition of a good answer. The teacher supplies it. And every proposed mark is reviewed and approved by the teacher before a student sees it.

Setting a paper. The teacher chooses the units, the structure, the balance of difficulty and the mark allocation, and the platform drafts a syllabus-aligned paper, with a teacher's version and a student's version. Questions can be edited, replaced or removed before it is used. A lesson planner works the same way, drafting objectives, activities and pacing from the curriculum content for the teacher to shape.


Where the teacher stays in the loop

For marking and teacher-created materials, the AI proposes. The teacher decides

MARKING HANDWRITTEN WORK



SETTING AN EXAM PAPER



A dashboard that says a pupil is struggling is a prompt to look, not a diagnosis. The same rule applies to a mark.

Seeing. The dashboard answers one question: which learners need my attention, and on which part of the lesson? It starts at the class, with who is active, which units are done, in progress or untouched, and a flag for learners whose pattern suggests they are struggling. It drills down to the individual learner and to the unit. At that level it shows three things side by side: how much of the unit is complete, the average score on its exercises, and the mastery estimate described above. They are not interchangeable. A learner can complete a unit with a middling score and a low estimate, and that combination is precisely the one a teacher wants to see.

A view for parents is coming, so that a parent can follow progress, build a revision paper, and ask the tutor how to help. The same rule applies to it as to everything above. A dashboard that says a pupil is struggling is a prompt to go and look, not a diagnosis, and the same is true of a mark.

A LEDGER PRINCIPLE

The AI drafts. The teacher decides. The time saved is for looking at children, not for looking at fewer of them.

Protect the learner

An adaptive tutor necessarily builds some state about the learner: an estimate of mastery, the misconceptions it has seen, a history of exercises, a schedule of reviews. That is what allows it to respond to the learner rather than to the question. It is also what makes a boundary essential. A mastery estimate exists to decide what support to offer next. It must not become a permanent judgement about the child.

Teachers need enough visibility to know where attention may be required. Parents need enough to support their child. Neither requires an unrestricted transcript of everything a fourteen-year-old has ever asked an AI tutor, and a learner who admits to the tutor that they have understood nothing since September needs to be able to say so without an audience.

The principle is simple to state. Keep what improves the learning decision. Expose it only to the people who need it for that purpose. And do not let a learning profile become a file on the person. The framework's [final era](#) imagines learning records that follow a person through life, and the paper it draws on is explicit that an early profile must never become an employment-screening instrument. The distance between a portable record of what I can do, that I control, and a file on me that started when I was fourteen is entirely a matter of design and governance. It has to be settled before the first student logs in, not after.

What we know, what we have tested, and what we still need to learn

None of these design choices began with one piece of research. They draw on a much longer tradition of mastery learning, formative assessment, intelligent tutoring systems, scaffolding and spaced practice, as well as the emerging evidence on generative AI covered in the previous essay. They were then reshaped by a real pilot, and by the ordinary work of building a product that teachers and students would use.

We have also been testing parts of the architecture ourselves. Roshan Summun, one of the architects behind myt Learn, used his MSc research at the University of Mauritius to examine one part of the problem more formally: whether a system that keeps an evolving estimate of understanding can make better progression decisions than one following fixed rules.

The useful result was not that the adaptive system “won”. It was that it exposed a trade-off. In simulation, it was substantially less likely to move a learner forward before they were ready. But it was also more likely to hold back someone who already understood. That is exactly the kind of trade-off an adaptive system has to make visible. Moving too quickly creates gaps. Being too cautious creates needless repetition.

But these were simulated learners, not Mauritian children in classrooms. The work tested whether the machinery behaved as intended. It did not test whether children learned more because of it. There is an obvious symmetry here. The previous essay argued that we should not mistake a good assisted answer for evidence of learning. We should apply the same discipline to ourselves. A successful simulation is not evidence of classroom learning. It tells us that the machinery behaves more like we intended. Real learners have to tell us whether that machinery actually helps.

Beyond the research, there are constraints that no design removes. These are the commitments we have chosen. They are not guarantees, and they do not remove the unresolved challenges of the AI-assisted era.

Availability is not use. The largest study in the previous essay, [two years across eighteen middle schools](#), found modest gains while substantive use of the tutor stayed limited. Every Grade 9 student having access to myt Learn would not mean every Grade 9 student using it in a way that could help. That is why our rollout model puts teachers at the centre of use inside the school day, and why the pilot's lesson mattered so much. Whether that is enough is one of the things we do not yet know.

The same tool is not the same opportunity. A learner with a quiet room, a device and an adult who asks how it went will get more from a tutor than a learner with none of those. Connectivity and devices are conditions we can influence. The adult and the quiet room are not. Support for participation belongs inside the model, not outside it as an explanation for failure.

Teachers' time is not automatically freed. A corrector that proposes marks still needs those marks reviewed. A generator still needs its paper edited. If the review step is honoured, as it must be, the potential saving is smaller than a product demonstration can make it look. We should not assume that introducing AI immediately reduces workload, and we do not.

The pathway ahead

I said at the start of this series that Era 3 is useful now and is not the endpoint. Having described one Era 3 product in detail, I hold both halves of that.

If this layer works, it will also start producing information we have never had at scale: where learners need more time, where they are ready to move faster, and where the curriculum's assumed sequence does not match the learner's actual one. A timetable cannot tell you that. A teacher responsible for thirty children cannot continuously record that level of detail across every concept and every learner. A system that records attempts and patterns of help across learners can begin to make it visible at a scale we have not had before. That is where Era 3 starts teaching us about Era 4. A school in which age organises belonging while readiness organises progression is not what we would be launching in 2027. It is what a launch may show us is possible, and where.

The architecture of mass schooling stays. Grade 9 students will still be Grade 9 students, in their classes, with their teachers, sitting their exam. What changes is that the learner who needed one more example can get it, while the learner who understood last week can keep moving rather than simply wait for the next class explanation. If we have built it right, they will be able to do more when the tutor is closed than they could before they opened it. That is the test.

WRITTEN BY

Veemal Gungadin

WITH ASSISTANCE FROM **Claude Fable 5.1** · Anthropic , **GPT-6 Astra** · OpenAI

Drafted and revised by the author. AI tools were used to sharpen structure and prose; conclusions are the author's own.

Published in *The Goodwill Ledger* at www.goodwill-led.com/essays/learning-before-answering/