15

Education in the AI era

Sample chapter from:
AI for Smarties
Understanding Artificial Intelligence
by Bertrand Meyer
Published by Smarties Press, July 2026
Available at https://ai-for-smarties.com

No area of human activity may be more affected by Modern-AI than education. You hear enthusiastic predictions about how AI will revolutionize education, and doomsday assessments that it is killing education. In the general spirit of the rest of this book, this short chapter puts aside both the hype and the gloom and tries to derive some practical lessons.

Often it seems that he enthusiasm comes from AI experts and others outside the educational system, and the doom from educators. Curious, is it not?

15.1 A REASONED APPROACH

This is a book about AI, with its fair share of enthusiasm (along with a clear description of the limitations and its doses of anti-hype vaccine). You might expect it to tell you to introduce AI everywhere in the curriculum.

Absolutely not. The injunctions to reinvent schools through AI are based on faith, not facts. The experiments that do exist yield mixed results at best. Hence the key advice:

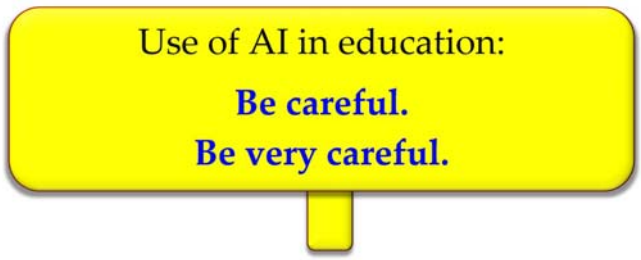

As noted in the previous chapter, education is an easy target for dogmatic promoters of drastic reforms, illustrating the H.L. Mencken that "*For every complex problem there is an answer that is clear, simple and wrong*". Education is a complex problem because education deals with one of the most delicate systems imaginable: the malleable, evolving mind of a young human. The current educational system is the result of 2500 years of evolution going back to the Greeks (more if we take non-European traditions into account). Its deficiencies are easy to see, but the replacements are much less obvious. Sweeping overhaul proposals are

only as good as their implementation. (For example, as discussed in 15.6 below, proposing new topics is easy; the hard question is what to remove.)

> The title of a 2025 French book [Alexandre 2025] translates loosely as *Stop Going to School: Learning Differently in the Era of AI*. Hmm.

"Stop your studies"

The main obstacle to making effective educational change is that credible empirical studies are rare, for a good reason: evaluating a teaching method involves treating students as guinea pigs, and if one approach turns out to work better than another, sacrificing the learning process for the students in the less fortunate group. You cannot do that.

Good educational research does proceed anyway, in particular research involving the use of AI, but at every step we must remember the principle stated above and avoid playing gratuitously with the educational process and young people's future.

15.2 EASY ANSWERS

How much AI should be included in the curriculum? Two opposite views are widely encountered:

- [ALL] AI is here to stay; we cannot ignore it in education any more than in any other field. Any attempt to resist the trend is going against progress and akin to putting our heads in the sand. Since we cannot fight it, we have to accommodate it at every stage of the educational system.
- [NONE] AI is a threat to education as we know it. Keep it out of the educational system.

Each actually has sound arguments in its favor, but applying either in full is unrealistic. We can no more let AI transform education unchecked than we can keep it away entirely. How do we find the right balance?

15.3 CHEATING

One of the biggest AI-related concerns worrying educators is the difficulty of giving and grading a take-home test today, because students — at all levels, even primary school — know that they can let an LLM solve it for them. For some teachers this problem has become a nightmare.

As a first observation, we note that AI offense can be met with AI defense. This approach does not solve the problem but mitigates it. The idea is to use AI tools to identify results inappropriate produced by AI tools.

AI offense, AI defense

Let me mention a recent personal experience. (An anecdote does not prove anything — although sometimes it can help *disprove* a thesis — but it can illustrate

a solution, which is the case here.) For a university course on an advanced topic (software verification) some students were remote and so was the final exam. I had explicitly disallowed the use of AI tools, for the simple reason that the course covered known topics, on which an LLM could be expected to produce many answers correctly, depriving the exam of any meaning. The rules were clear and the students had signed an honor code. For practical reasons we were not using a proctoring tool (more on those below).

Most of the students turned in a normal set of answers but I could immediately see that one of them had produced all answers through an LLM. They were not only correct but impeccable in form; in addition, instead of just addressing the questions they had profuse details, making them more like little tutorials. Plus, they included the kind of thing one *never* sees in a student exam response written under strict time limits: bullet points, boldface…

My certainty that the student had cheated would have no formal value if I filed a complaint to the academic administration; students have considerable protection, and an instructor's personal conviction carries no weight.

I could, however, use a firmer approach. I fed in one of the answers to an LLM, asking for a probability that the answer was AI-generated, and received an estimate of 65%, backed by precise arguments. I added another one of the students' answers; the estimate went up to 87%. Third answer — now 98%, with a detailed, carefully constructed argument (comparable to the analysis, in the first example of our introductory chapter, "An invitation one cannot refuse?", 2.1, page 11, of why an email was a scam).

That argument carried the day. I sent the full LLM interaction to the student, who immediately confessed to the cheating and apologized, avoiding an escalation and enabling a no-drama resolution. Once the fraud had been uncovered using the same kind of tool that had produced it in the first place, it was clear that denial would be impossible if escalation did happen.

Students who are considering taking the easy way should think further and realize that the game has two players: instructors have access not only to all the tools students can use but to others as well, such as sophisticated AI-based plagiarism checkers.

> Caveat: I mentioned that this example is an anecdote. Running a suspicious exam through an AI cheating detection tool is no panacea. As emphasized throughout this book, Modern-AI tools are probabilistic. However high the probability of correctness, it will not be 1: if we run such tools on millions of answers, inevitably a few thousand will be wrongly flagged. Such cases of unjustly accused students have already occurred. For any serious decision, human judgment is essential.

Monitoring tools

The heavy-duty way to combat cheating, for online exams, is to use one of the various proctoring systems — using AI technology of course — available on the market. They monitor the students' webcam, triggering an alarm if they move too much or another person enters the frame, but also track all their computer activities including Web access and other Internet traffic. This solution pleases academic management because of its promise of automation and exhaustiveness. It solves the problem in principle, but requires some caution:

- Practice is seldom as smooth as theory. The local system administrator cannot fully control a setup which involves remote installations with their own platforms, web browsers, Internet connections, webcam and microphones. There are many points of failure, not necessarily the students' fault. There is a risk that if any element in this delicate chain breaks one has to call off the whole mechanism and (paradoxically) revert to a insecure mechanism. Even if the setup finally works, precious time is often wasted at the beginning of the exam, leading to time extensions and other complications.
- Getting things right puts pressure on everyone. An exam is by nature nerve-wracking; attention devoted to the machinery of the exam rather than its substance has no educational value.
- The tools are expensive, not only financially but because they take time to operate, at the expense of time devoted to educational goals.
- While cheating is a serious problem, in particular because it harms honest students, it is not necessarily ideal to turn exams into police operations.

For those reasons, such proctoring tools are best reserved for high-stakes situation in which remote exams are the only solution and the decisions (such as admission into a school or program) are critical. For more normal course assessment, they can be overkill.

How do we proceed in such cases? No one has the answer yet. The best we can do is to put in place strong regulations, explain clearly and repeatedly what is permitted and prohibited, use honor codes with serious consequences for violations, and get all the help we can from technology, with due precaution.

15.4 WHAT HAPPENS TO HOMEWORK?

For centuries, the educational system has relied, along with exams and other in-class tests, on homework. You give an assignment that the student takes home, turning an answer a few days later. While some amount of dubious practices (not all necessarily amounting to cheating) have always been possible, for example by consorting with other students beyond what is officially permitted or getting help from family, homework has been a widely applied educational

technique, focusing on in-depth work such as solving complex problems or writing essays, with the benefit of access to books and other resources, and more time than students typically have in a classroom.

What is the role of homework in today's AI-rich tech scene? No beating around the bush:

Unpleasant as this truth may sound, we can no longer rely on the graded homework, a time-honored technique as a key element of our educational panoply. The reason is simple: for most courses, particularly in primary and secondary school but also in undergraduate education and most graduate courses, it is just to easy to feed the homework into an LLM and get a good answer.

The instructor may disallow the practice and apply the kind of technique illustrated by the anecdote at the beginning of this chapter. This approach does not scale up. The remote-exam setup of that example was specific, and justifies using technology to track the occasional cheater. For homeworks, the difference in effort between solving the problem yourself and asking an LLM is so huge — at least an order of magnitude — as to make resisting the temptation heroic. Most students are not heroes. Few teachers will want to devote the rest of their career to playing the cat-and-mouse game (where both the cat and the mouse are AI tools).

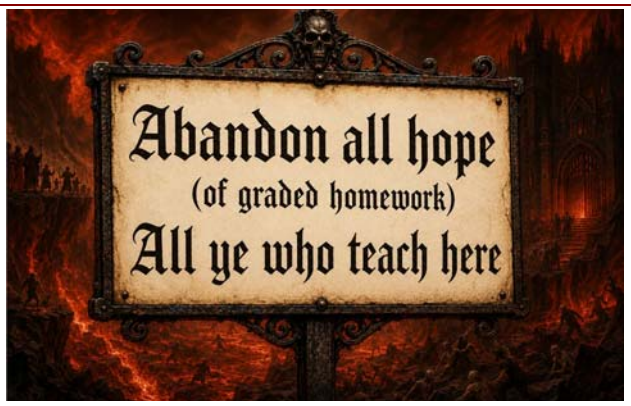

The statement applies to *graded* homework. You can still use homework as a way to make students perform some individual work, if you convince them that it will be in their benefit — even allowing them to use AI tools as a help towards the educational goal, not to produce the result. Getting closer to self-learning as

discussed in the preceding chapter, this approach assumes students have enough maturity to take the assignment seriously for their own benefit.

Homework may retain its value as a teaching instrument, but loses its value as an assessment instrument.

15.5 USING AI IN TEACHING

The preceding observations lead to the general question of how much AI to allow in the teaching process. The general injunction is just as clear here:

This warning is particularly applicable to primary and secondary ("K-12") education, but remains relevant for a large part of the university-level curriculum. It covers the teaching of topics other than AI itself (say history or mathematics, where AI would be a tool). Teaching of AI itself is the topic of the next section.

I realize that the advice will shock some readers, especially those who (perhaps with the help of the preceding chapters) have become enthusiastic about AI. But it is incontrovertible. It applies in particular to LLMs, which in the context of teaching provide the kind of shortcut that we precisely want to avoid: shortcuts to thinking. It suffers some exceptions:

- The more advanced the teaching level, the more flexible you can be, particularly if you trust the students' drive and sense of responsibility. It is particularly at the primary, secondary and undergraduate levels that we need to close off the unnecessary shortcuts.
- Some AI tools are part of the toolset used by practitioners of the subject being taught. Astronomers will be introduced to tools analyzing snapshots of space. This case tends to happen mostly in advanced courses.
- More generally, students will obviously use AI in courses about AI and, more generally, computer science topics. See the next section.

Outside of such cases, AI tools, particularly LLMs are **an obstacle to teaching, not a help**. The fundamental observation is simple: education at the K-12 and early university levels teaches *topics*, but it most importantly teaches *thinking* — including *learning*. It teaches the art of analysis and synthesis, developing a critical mindset, strategies for acquiring new knowledge, structuring acquired knowledge, reformulating in your own terms what you have learned, explaining it to others, knowing what to learn next, organizing your work. These are

precisely the tasks that AI tools perform for you. If you use them in your learning process, you forsake the benefits.

If you do possess those basic skills, then AI tools function as an **intelligence amplifier**; the previous chapter has described an effective process to achieve the amplification. But for a not yet mature student they are a catastrophe since they function as an **intelligence substitute**.

As already noted, current AI technology is still too young to have been subjected to definitive studies of its effectiveness in teaching environments. The studies that do exist all point in the same direction:

- Letting students use AI tools has a good short-term effect. Students appear to get the hang of new topics quickly.
- After a while, the picture changes. Because the students skipped the steps of struggling with a topic and forming their own internalized version of it, adapted to their own specific mindset — a sometimes painful but eventually productive process — they do not master it in depth, and are subject to forgetting quickly whatever they have learned.

"No pain, no gain" is unpleasant to hear but seems to be a fundamental component of learning.


Elementary school teachers picket against use of calculators in grade school

Math teachers protest against calculator use

Unconditional defenders of AI in schools can dismiss these arguments as reactionary and point to previous cases of pointless resistance to the introduction of technology into classrooms, as when US teachers picketed against calculators in 1966. In the 1970s, some computer science professors resisted the use of interactive programming environments, because students could now become sloppy (whereas being forced to turn in a single "batch job" every few hours, and to start again if there was a single mistake, forced them to be careful).

Such anecdotal precedents do not affect the argument. Any good tool frees humans from some tasks. In an educational setting, the majority of students will lose the corresponding skills. With the spread of calculators, fewer people are good at mental calculation today. Programmers can skip learning details of programming language syntax (the tools will fix them). The question is whether such skill loss is acceptable:

- Society can live with most students not mastering mental arithmetic or programming language syntax details. We actually prefer that they focus on learning more advanced skills.
- AI tools, however, substitute for *intelligent thinking*! That is not a skill we should be prepared to sacrifice teaching.

If AI is to be your intelligence amplifier rather than intelligence substitute, it can only amplify what you already have (in some quantity). At school, we are teaching intelligence, and do not need shortcuts.

The practical lesson for teachers at elementary level is clear, captured by the principle at the beginning of this section: AI is a powerful tool if you use it as a teacher, but as a tool for your students it is a competitor, often standing in the way of your educational goals. This point is worth emphasizing again:

This injunction applies to today's technology. It is possible that tools will come around that are well integrated into the teaching process and supportive of it. Until then, caution should guide us.

Keeping AI out of the classroom can be difficult, but is not impossible. After all many schools have realized that smartphones in the classroom bring far more trouble than benefits and have succeed in banning them. (It has been widely reported that the most famous titans of the tech industry, when it comes to their own children in their own home, ban technology.)

There remains a role, limited but crucial, for AI in the classroom: treating AI not as a tool to learn other subjects but as a subject in itself. We turn to it now.

15.6 STUDYING AI

Excluding AI as a learning tool is not a Luddite reaction but results from an analysis of what school is about. We cannot let students rely on tools to skip the very goal of going to school. Part of that goal, however, is to learn to function well in society, and AI is certain to play a key role in society; every student should be ready to use AI well.

This objective is essential if we want to graduate students who will take advantage of AI (as their intelligence amplifier) — and not be taken advantage of by AI and by an elite which masters AI.

As a consequence, school programs should include artificial intelligence as a subject of study. The nature of what should be taught about AI varies with the educational level (this is not the place to make detailed curriculum proposals), but here are some of the elements — similar in many cases to the topics of the present book — that can lay a claim to a spot in the curriculum:

- What AI is and is not (demystification).
- What AI can and cannot do.

- Using AI responsibly (introduction to AI ethics).
- Using AI for problem solving; prompt engineering.
- Mathematical basis for AI (see chapter 8): linear algebra, differential and integral calculus, optimization methods.
- Fundamental mechanisms of Modern-AI: training, training sets, neural networks, forward and backward propagation — and at more advanced levels optimization, gradient descent etc. (see chapters 7 and 9 to 11).
- Computer science for AI: algorithms, data structures, programming languages, requirements engineering, software verification.

An important problem is the fixed-pied nature of educational programs. As anyone who has ever participated in curriculum design at any level (primary school, secondary school, university) knows, the game is tricky; getting everyone to agree to a new topic ("*Schools should be teaching X!*") is possible, but the next question is "*What topics Y do we remove to make place for X?*", and enough people will protest vocally enough as soon as you put forward a concrete *Y*.

No such proposal will appear here, but we may note that some of the more advanced topics are candidates for integration into courses in mathematics and (as increasingly present in school curricula) computer science.

While the implementation may be delicate, the principle is clear if we want the citizenry of the future to master AI and not fall prey to it:

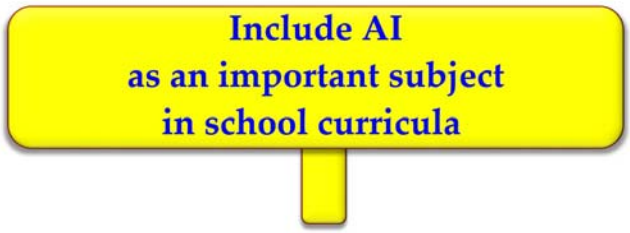

15.X Exercises for chapter 15

15.X1 Extreme approaches to artificial intelligence in education

- A. Assume a primary school implements a policy of complete, unrestricted AI integration across all subjects. Analyze specific practical consequences of this decision on the day-to-day operations of a classroom.
- B. Assume that the policy is the reverse one: complete exclusion of AI tools. Can you think of a concrete scenario in which the policy fails due to external societal factors?

Other exercises for this chapter appear at ai-for-smarties.com/exercises.