1. AI in Education

- 1.1 Potentials & Risks
- 1.2 Teaching About vs Teaching with AI
- 1.3 Responsible Use of AI
- 1.4 Compliant Use of AI
- 1.5 Current State of AI in Education
- 1.6 Legislation

1.1 Potentials & Risks

> *If you remember only one sentence form this page:* AI has made its way into schools like a medicine that comes without a patient information leaflet. As a teacher, you don’t know when to use AI, you lack dosage instructions, and there is no information on the side effects. This chapter serves as the patient information leaflet for the Playbook.

Should AI come with a leaflet like a medicine?

You know that slip of paper tucked inside every medicine box? The one you likely toss aside without a glance? It may seem dull, but this little paper is a quiet triumph of the last century. That leaflet spells out four essentials: what the medicine treats, how much to take, the surprises it might bring, and who should steer clear.

Behind every line on that leaflet stands a whole machinery: clinical trials, regulators, and a responsibility to reveal mistakes. No one believes the leaflet itself heals you. Its real power is making the medicine safe for anyone. That is the magic we are missing when we let AI teach without guidance.

The AI tools now quietly settling into your school inherit a problem that is older than AI: educational technology has rarely come with instructions. In the UK, a review found that only 7 out of 100 education technology companies had put their products through a randomised trial, and only 12 out of 100 had sought outside certification. Schools rarely ask for more. In a survey across 17 American states, only 11 out of 100 teachers and administrators had requested peer-reviewed evidence before adopting a new tool.

The result, in many cases, is wasted money. In the US, on average two-thirds of education software licences went unused, and 98 out of 100 were never used intensively. The EdTech Genome Project, also in the US, looked at about 7,000 teaching tools worth 13 billion dollars and found that 85 out of 100 were either a poor fit or implemented incorrectly.

This chapter may look like the dry, necessary section, but it reveals the hidden stories behind the rules. It tells what really happened, who was affected, and shines a light on the side effects, even those that quietly reached children who never signed up.

Potentials of AI

Most studies on AI's potential focus on university students. School pupils have so far featured relatively rarely in these studies. And most studies relate to short-term interventions lasting from a few hours to a few weeks. Consequently, the so-called novelty effect may have positively influenced the results. A large review of 68 experiments found exactly this pattern. The shorter the study and the smaller the number of particpants, the larger was the reported benefit of AI.

(1) Let AI explain a new topic and keep the practice for your lesson

Imagine this. Tomorrow, when a pupil misses a topic or needs that tricky concept explained yet again, let an AI tutor handle the first explanation. Reserve the practice for your lesson, where real learning comes alive.

Why this matters. In a well-controlled trial, an AI tutor taught new material at least as well as a good lesson. In a Harvard study, 194 physics students learned two topics. For one topic, they used an AI tutor developed for this purpose. For the other topic, they learned in a well-run active-learning class – with discussion and problem-solving in pairs, not a lecture. Each student completed both approaches, so no one can claim that the stronger group received the better treatment. In a test conducted immediately afterwards, the students taught by the AI achieved significantly higher results. The researchers recorded a median time of 49 minutes to complete the task, compared with the 60 minutes allocated for the lesson in the timetable. The authors are cautious in their interpretation of these results: the material was new to the students, and much of the task simply involved explaining it well. They explicitly decline to extend their result to tasks that require pulling several ideas together. And these were undergraduate students, not eighth-graders.

A large review pooling the results of 228 studies points the same way, with one important limit. Generative AI produced larger gains in knowing and understanding than older forms of AI such as intelligent tutoring systems and adaptive practice software. But this was the only one of the four outcomes the review measured where generative AI had the edge.

Go deeper. 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,* Article 17458. Free access, DOI 10.1038/s41598-025-97652-6. The article describes how the tutor was designed and controlled. This is useful for anyone wishing to develop an AI tutor themselves.

(2) Making better use of lesson planning time

Imagine this. Save some of the time you need for lesson planning by using AI to draft your first lesson plan and then revise it thoroughly.

Why this matters. No one disputes that teachers' workloads are heavy. Therefore, one potential benefit of AI is using working hours more efficiently. A trial in England investigated whether targeted use of AI for lesson planning can reduce teachers' workload. Researchers randomly assigned 259 science teachers at 68 schools either to a group that used ChatGPT together with an online guide to using it well, or to a group that was asked not to use generative AI at all. **Teachers in the ChatGPT group spent about 56 minutes a week planning science lessons for Years 7 and 8 (ages 11 to 13), compared with about 81 minutes in the comparison group.** A panel of subject experts assessed the quality of the lesson plans. They did not know whether the plans had been created with or without the aid of AI. The experts found no evidence to suggest that the quality of the lesson plans differed. However, two points must be highlighted. (1) **The teachers used the tool less and less as the weeks went on.** The evaluators read this as teachers moving from a learning phase into routine; it may also mean that the saving does not sustain itself without support. (2) The trial measured the time spent by teachers on lesson planning, not their pupils' learning outcomes.

Voices from the field

"Everything comes back to time. Even teachers who see the value of AI rarely have the time learn how to use it."

Go deeper. Roy, P., Poet, H., Staunton, R., Aston, K., & Thomas, D. (2024, 12 December). *ChatGPT in lesson preparation: A Teacher Choices Trial.* NFER, commissioned by the Education Endowment Foundation and the Hg Foundation. Free at educationendowmentfoundation.org.uk.

(3) Providing better support for learners who do not benefit fully from standard lessons

Imagine this. Your school decides how it wishes to spend its funds. It wants to purchase an AI chatbot. How likely do you think it is that the chatbot providers will provide you with evidence of the bot's benefits for different subjects?

Why this matters. The provider may point to research on learners with disabilities – one of the longest-studied fields of AI in education, and one that is easily overlooked when funds are allocated. In an analysis of 29 studies involving 41 groups of learners with disabilities, researchers found moderate benefits across a wide range of settings. But read the small print. In **seven out of ten of these studies**, 'AI' meant a **robot**, usually a humanoid one, often steered by a researcher out of the child's sight. Only one in five studies used software. So this is no proof that a chatbot will help an autistic pupil in your class. Nor could the reviewers say **what** makes these interventions effective. Based on their research, they could not identify why the studies that worked differed from those that did not.

Go deeper. Zhang, L., Carter, R. A., Jr., Liu, Y., & Peng, P. (2026). Let's CHAT about artificial intelligence for students with disabilities: A systematic literature review and meta-analysis. *Review of Educational Research, 96*(1), 215–257. Free, DOI 10.3102/00346543241293424. The real gem is

the table of included studies: it instantly reveals just how little of this field actually focuses on the tools currently being marketed to schools.

(4) Teaching learners how AI fails

Imagine this. Ask your class to create ten images of 'a doctor' and ten of 'a nurse.' Tally up who appears in each set, then spark a discussion: where did these patterns come from, and who made those choices?



Fig. 1: One shot prompt: a doctor, a nurse (Image generated with NanoBanana 2)

Why this matters. The classic version of this lesson feels outdated, and that is exactly why it deserves attention. For years, the demonstration was predictable: ask for a doctor, and a man appears. In a test by the non-profit AlgorithmWatch, Midjourney almost always pictured doctors as white men, while Adobe Firefly made visible, if imperfect, attempts to diversify. A peer-reviewed study of 2,400 AI-generated images of surgeons found the same split between products: two of three popular generators showed almost no women, while the third came close to the real proportion of women in surgery.

Try the same prompt today, and you might see a female doctor or a male nurse appear (Fig. 1). But do not be fooled into thinking the bias has vanished. A 2026 analysis of 1,344 images from three popular generators revealed the opposite: when asked for a 'competent person,' one system

showed no women.

What changed is not the bias but the layer it lives in – and that is the important lesson. Human decisions stand behind any of these pictures: what the system was trained on, and what its makers afterwards decided it should show. The second is adjustable, invisible, and made by a company.

The first layer of human decisions is what the system has absorbed. Millions of images, crafted and labelled by people, each one echoing the world as it already was. This is the layer that the classic bias lesson reveals, and pupils grasp it easily: the machine became a reflection of what it was fed.

The second layer is shaped by what a company chooses to show you next, yet this influence leaves no trace in the image itself. Three things unfold here, each one described by the creators in their own words.

- **Your words are rewritten before the model receives them.** From OpenAI's own developer documentation: "we use GPT-4 to optimize all of your prompts before they're passed to DALL-E" — and "this feature isn't able to be disabled at the moment." What your pupil typed is not what the machine was asked.
- **Appearance gets filled in when the prompt leaves it open.** OpenAI, describing its diversity technique: "This technique is applied at the system level when DALL-E is given a prompt describing a person that does not specify race or gender, like 'firefighter.'" Afterwards, users were twelve times more likely to say the images showed people of diverse backgrounds. "A doctor" is exactly that kind of prompt.
- **And that tuning can miss.** Google, after pausing Gemini's image generation in 2024: "Our tuning to ensure that Gemini showed a range of people failed to account for cases that should clearly not show a range."

So when a pupil types "a doctor" and a woman appears, the image offers no clue about which layer brought her into being. **That uncertainty is not a flaw in the lesson. It is the lesson itself.** The old story claimed the machine simply soaked up society's prejudices. The new story cuts deeper and lasts longer: A company made a choice! You cannot spot that choice in the outcome, and it might shift on a Tuesday, quietly, without warning. That truth endures through every new model. The stereotype lesson fades away.

This is the lesson that deserves a student's attention, and it will outlast the next model update: **the image is not a simple reflection of reality, and it did not happen by chance. Someone made a choice.**

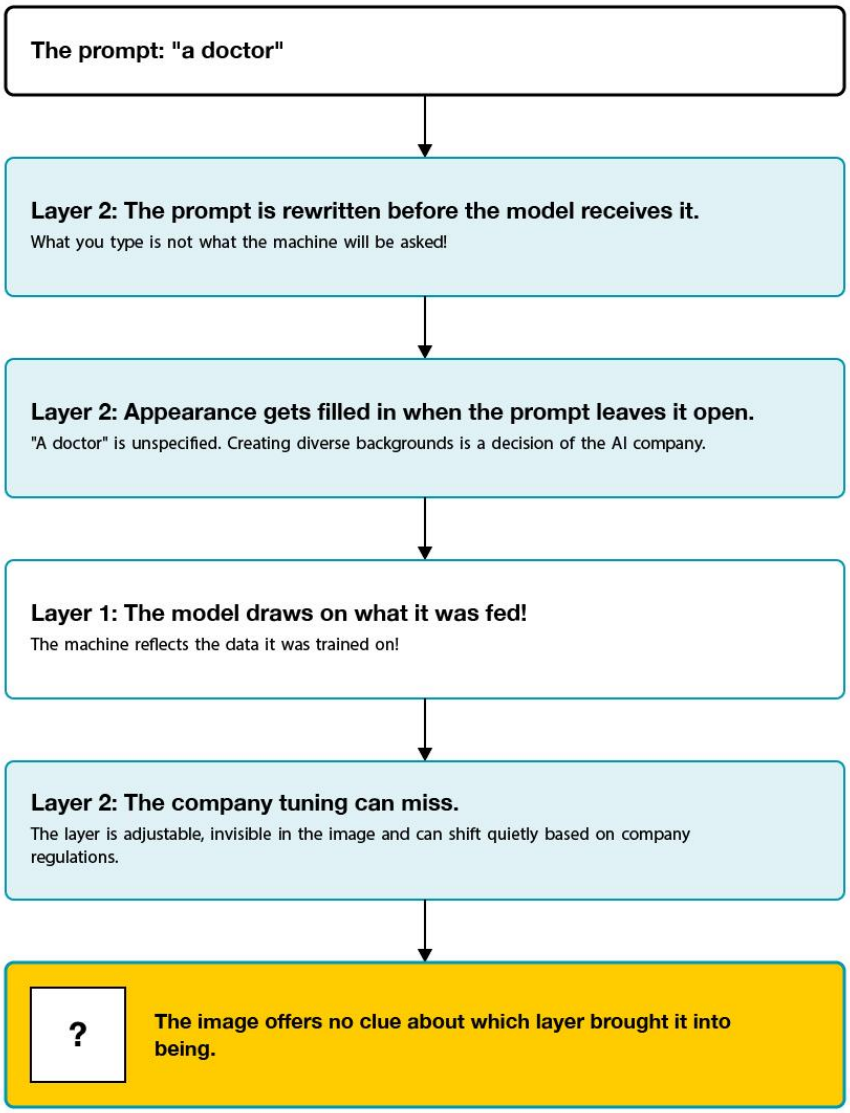


Fig. 2: Why the picture based on a prompt like "draw a doctor" looks like it does and why the picture itself won't tell you (Image generated with NanoBanana 2)

This lesson can start early. When 209 Finnish students in 12 classes, grades 6 and 7, explored this topic, the number who could explain the bias using data jumped from about 7 out of 100 to 4 out of 100. The honest reservation: the study measured what children could explain, not what they would do later.

Voices from the field

"Bias is not unique to AI. Even textbooks can contain one. What matters is knowing about chance, about the data behind an answer & about possible distortions."

Go deeper. Vartiainen, H., Kahila, J., Tedre, M., López-Pernas, S., & Pope, N. (2025). Enhancing children's understanding of algorithmic biases in and with text-to-image generative AI. *New Media & Society, 27*(9), 5342–5368. Free, DOI 10.1177/14614448241252820. The study had no

comparison group – read it as a promising classroom design, not as proof.

Weinmann, H., Messingschlager, T. V., & Appel, M. (2026). Gender bias in text-to-image generative artificial intelligence: Neglect and stereotypical presentations across three popular platforms. *New Media & Society*. Online first. Free, DOI 10.1177/14614448261435197.

Ali, R., et al. (2024). Demographic representation in 3 leading artificial intelligence text-to-image generators. *JAMA Surgery, 159*(1), 87–95. DOI 10.1001/jamasurg.2023.5695.

The vendors' own statements: OpenAI (2022), *Reducing bias and improving safety in DALL·E 2*; Depue, W. (2023), *What's new with DALL·E 3?*, OpenAI Cookbook; Raghavan, P. (2024), *Gemini image generation got it wrong. We'll do better*, Google.

How solid are the research results on the potential of AI at this very moment?

The most important thing. If someone quotes a study to you, ask two questions before agreeing: "Who were the learners, and how long did the study last?"

Why this matters. In 2025, a meta-analysis of 51 studies reported a **large benefit** of ChatGPT for students' academic performance. The article was read hundreds of thousands of times and served as the basis for numerous confident statements in staff rooms. **On 22 April 2026, the journal retracted the article.** A *retraction* means that a journal officially withdraws something it has published; the article remains visible, marked with the note 'RETRACTED', so that anyone who has cited it can recognise this. Two researchers who had not been involved in the study had identified discrepancies in the analysis. The journal's editor wrote that the problems 'undermine the editor's confidence in the validity of the analysis'. The authors did not respond to correspondence on this matter.

This does not mean that AI has no impact. There are reputable studies that demonstrate moderate benefits. This field of research is also not yet well defined because it is relatively new and highly dynamic. It means that the figure that was frequently cited was incorrect, and that it took a year and two external experts to realise this.

Another reservation is worth mentioning. In mathematics, at least, the closer the evidence gets to school, the smaller the effect. Across subjects the picture is mixed: some recent reviews find larger effects in schools than at universities – but the school studies are still few and small.

Go deeper. The retraction itself is one page, free, and worth reading precisely because it is short: Wang, J., & Fan, W. (2026). Retraction note: The effect of ChatGPT on students' learning performance, learning perception, and higher-order thinking. *Humanities and Social Sciences Communications, 13*, 528. DOI 10.1057/s41599-026-07310-z.

The **retracted** original was *HSSC 12,* 621 (2025). Still standing: Wu, X., Zhu, P., Zhang, J., Yin, M., & Wang, Y. (2026). ChatGPT's impact on student learning outcomes: A meta-analysis of 35 experimental studies. *Humanities and Social Sciences Communications, 13*, 684. Free, DOI

10.1057/s41599-026-07019-z.

The school-only figure for mathematics: Yi, L., Liu, D., Jiang, T., & Xian, Y. (2025). The effectiveness of AI on K-12 students' mathematics learning: A systematic review and meta-analysis. *International Journal of Science and Mathematics Education, 23*(4), 1105–1126. DOI 10.1007/s10763-024-10499-7.

A broader, more recent pooled analysis: Han, X., Peng, H., & Liu, M. (2025). The impact of GenAI on learning outcomes: A systematic review and meta-analysis of experimental studies. *Educational Research Review, 48*, 100714. DOI 10.1016/j.edurev.2025.100714.

Risks of AI

(1) Doing something is not the same as learning

Imagine this. Before you set a homework task a chatbot can complete, think about the purpose of the task. Might it be a good idea to change the task?

Why this matters. Do you learn faster with a chatbot? Around a thousand pupils in Years 9 to 11 practised maths in three groups: one using a standard chatbot, one using a version programmed to withhold the answer and ask a follow-up question, and one without a chatbot. Whilst using the chatbot, the group with the standard chatbot completed 48 per cent more practice exercises than the group without one. Afterwards, everyone sat the same exam without a chatbot. The result, which seemed surprising at first, was that the group using the standard chatbot performed 17 per cent worse than the pupils who had not used a chatbot. The group that had worked with the chatbot which withheld the answer showed no drop in performance in the test compared to the group without a chatbot, but practised more. The pupils using the standard chatbot were neither lazy nor did they cheat. They worked harder and learnt less, because what the tool eliminated was not the effort, but the cognitive challenge that constitutes learning.

Voices from the field
“AI makes it very easy to get a solution without understanding it.”

Go deeper. Bastani, H., Bastani, O., S ngu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *PNAS, 122*, e2422633122. DOI 10.1073/pnas.24 2633122.

(2) AI Detectors punish the wrong pupils

Imagine this. You receive a student’s essay to grade. Curious about its authenticity, you feed it into an AI detector. The result flashes: the essay was almost certainly written by AI.

Why it matters. You use the AI detector to ensure fairness, **yet for a particular group, the detectors have exactly the opposite effect.** A detector does not recognise machine-generated text. It measures how predictable the wording is. And that is precisely where the problem lies, for example, when a pupil is writing in a foreign language, is unpractised in writing, or has a limited everyday vocabulary. Their writing is predictable. When seven widely used detectors analysed 91 essays written for the TOEFL English test by non-native speakers – collected from a Chinese online forum – they wrongly flagged them as machine-generated in around six out of ten cases. Essays by American eighth-graders writing in their native English were almost all classified correctly. A tool that triggers false positives for exactly one group is not neutral. Furthermore, this group is the least able to defend itself.

Go deeper. Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. *Patterns, 4*(7), 100779. Free, DOI 10.1016/j.patter.2023.100779.

(3) Where on earth will the pupils' data be?

Imagine this. You use an AI tool your school offers. Try to find out where the pupils' data go.

Why it matters. Can your school ensure that the data is not passed on to a third party? A third party is a company that is neither your school nor the app’s developer. Typically, these are advertising or analytics companies that receive data whilst the pupil is working. No one at the school chose this company. When researchers carried out a technical analysis of 163 learning products recommended by governments during the pandemic-related school closures, 145 of them processed children’s data in a way that jeopardised or infringed their rights by transmitting data to 196 third-party companies or granting them access, predominantly from the advertising industry. **Of the 42 governments that developed their own products rather than purchasing them, 39 had developed products with the same problem. The study covers products from 2021, and some may have changed since then**. However, we cannot rule out that the same applies to AI applications.

Go deeper. Human Rights Watch (2022). "*How dare they peep into my private life?:* Children's rights violations by governments that endorsed online learning during the Covid-19 pandemic. Free at hrw.org. The country annexes let you look up what was endorsed where.

(4) AI arrives last where it is needed most

Imagine this. How much training did you already receive on learning about AI or subject-specific teaching with AI?

Why this matters. In one experiment, adults with and without a university education solved a work-style business task. With an AI assistant, about three-quarters of the gap between the two groups disappeared. When the researchers then took the assistant away, a sizeable part of the gap came back. The authors conclude that the AI did much of the work, and that only part of the gain survived once it was gone.

A survey of American school principals adds the second half of the picture: the more disadvantaged a school's pupils, the less the school had done to integrate AI – fewer policies, less teacher training, less guidance on pupil use. That is an association, not proof of cause.

Taken together: a good tool can narrow a gap while it is in the learner's hands. But it tends to reach the schools that need it least, and only part of its benefit stays behind when it is removed.

German figures show the same pattern. Among 1,590 young people aged 14 to 20, 80 per cent from the most affluent families but only 55 per cent from the least affluent saw AI as an opportunity for school. Usage follows the same gradient: among young people who use AI tools, 70 per cent of Gymnasium pupils use them for school or homework, compared with 58 per cent at other secondary schools.

What about teacher professional development? A survey of more than ten thousand teachers in England found that 45 per cent of teachers at private schools had already received formal training on AI, compared with 21 per cent at state schools. Within the state system, the figure ranged from 26 per cent at schools with the most affluent intakes to 18 per cent at those with the least affluent.

Go deeper. Cruces, G., Fernández Meijide, D., Galiani, S., Gálvez, R. H., & Lombardi, M. (2026). *Does generative AI narrow education-based productivity gaps? Evidence from a randomized experiment* (NBER Working Paper 34851). A working paper, not yet peer-reviewed. Campos, C., & Singleton, J. D. (2026). *AI diffusion gaps: Unequal integration of AI across K-12 schools* (NBER Working Paper 35347). Latham, K., & Montacute, R. (2025). *Artificial advantage? AI in the classroom and the inequality gap*. The Sutton Trust. Free at suttontrust.com. Vodafone Stiftung Deutschland (2024), *Pioniere des Wandels*; mpfs (2024), *JIM-Studie 2024*. Both free, in German.

(5) What improves when learning with AI is not necessarily what you see

Imagine this. You introduce a new chatbot into your teaching and the pupils are enthusiastic about using it.

Why this matters. A large review of 228 studies looked at four different things AI might improve: knowing and understanding; applying what one knows; the affective side, such as motivation and engagement; and the ability to plan, monitor and correct one's own learning. Three of the four improved. The last one did not, as far as the studies could tell.

- **For knowing and understanding,** generative AI did better than older forms of AI such as intelligent tutoring systems and adaptive practice software.
- **For applying knowledge, and for the affective side,** generative AI did no better than the older software many schools already own.
- **For self-steering,** the pooled effect was small and could have been chance – and there were too few studies to ask which kind of AI might help at all.

Only a handful of the results came from randomised trials, and the enthusiasm you see in the first weeks is exactly what short studies capture best.

Here’s a quick side note: **in this review, the greatest gains were found in the arts and humanities.**

Go deeper. Yeo, G. H., & Lansford, J. E. (2025). Effects of artificial intelligence on educational functioning: A review and meta-analysis. *Educational Psychology Review, 37*(4), article 110. DOI 10.1007/s10648-025-10085-5.

Go deeper. Yeo, G. H. & Lansford, J. E. (2025). Effects of artificial intelligence on educational functioning: A review and meta-analysis. *Educational Psychology Review, 37*(4), article 110. DOI 10.1007/s10648-025-10085-5.

1.2 Teaching About vs Teaching with AI

Take Home message

- ***AI is a satnav.*** AI can save the driver who already carries the map and prevent the one who doesn't from learning. The effect depends on the person, not the product.
- ***We argue that teaching about AI is the precondition for teaching with AI.*** You need to know the machine to check the answer it produced.
- ***Knowing is not doing.*** Teachers rate their AI knowledge above their ability to teach with it, and training alone does not close that gap.
- ***Teachers use AI backstage, not in class.*** In Germany, half of teachers use AI several times a month or more; of those, fewer than a third let pupils work with it in class.
- ***Pedagogy beats technology.*** Using AI to rework and build on one's own work beats using it only to practise or look things up; working in groups beats working alone; feedback beats gamification.
- ***Better feedback is not better learning.*** In one experiment, AI comments judged better than a teacher's led to revisions that were no better. What decides is whether feedback is *taken up.*
- ***AI raises output, not self-directed learning.***
- ***Clear rules and heavy use appear together.*** In a large US survey, schools with an AI policy had more teachers using AI – and so saved more time overall, not less.
- ***Guard your own map.*** *Relief becomes deskilling where AI absorbs the work that builds and maintains judgement.*

> *If your remember only one sentence from this page:* **AI is a satnav. It saves the driver who already has a map in her head and quietly stops the driver who hasn't from ever building one**

The satnav problem

To drive a London taxi, you must first pass a knowledge test: 320 set routes across the 113 square miles within six miles of Charing Cross, for which Transport for London allows up to two years before the examinations even begin.

Examined drivers also use a satnav, and nobody thinks less of them for it. The knowledge test did not become worthless when the device arrived. The satnav became the thing that makes the drive safe. A driver with the city in her head uses the satnav for what she cannot know: the accident on the bridge or tonight's traffic. And she notices when it routes her down a street that has been dug up since April.

Now put a newcomer in the same taxi with the same satnav. She arrives at every address. For a while, she is indistinguishable from the experienced driver. But following the satnav’s instructions, she is building nothing. After a year of turning left when instructed, she still cannot cross the city without the device, and she cannot tell when it is wrong.

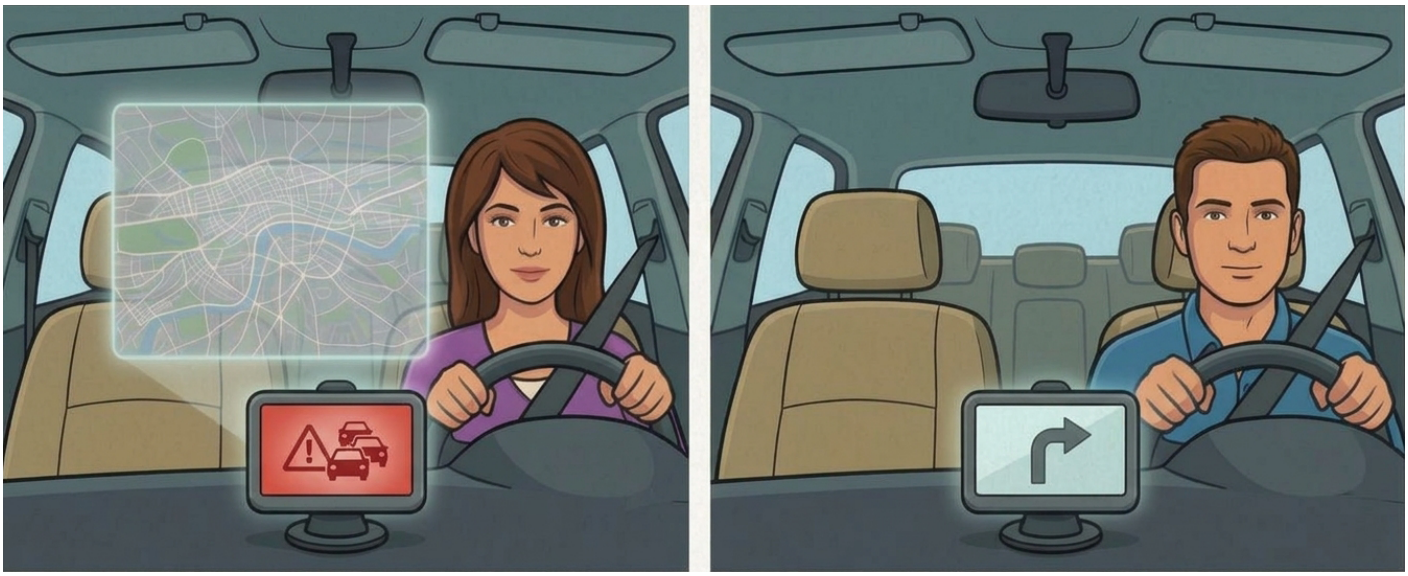


Fig. 1: The outcome depends on the driver. On the left side the driving begins with a map in the head. The the satnav adds what the driver cannot know. This is called "beneficial offloading". On the right the driver has no map in the head. The device is the only orientation for the driver. This is called "detrimental offloading" (Image generated with GoogleGemini).

Same satnav, two drivers, opposite outcomes. **The difference was never in the device.** We talk about more than an image. Studies of drivers find that people who rely heavily on satnav form a poorer mental map of their surroundings. A study over three years showed that heavier use went with a steeper decline in spatial memory. And psychologists have known for two decades that the same support which helps a beginner can get in the way of someone who already knows the subject. The same happens the other way around. The effect of an AI tool depends far more on what the user already knows than on the tool itself.

This is what researchers call a schema, the organised prior knowledge of a person. Where it is in place, handing work over to a machine frees you up. Researchers call this beneficial cognitive offloading. However, when knowledge needs to be developed, the same process can prevent it from forming. This is called detrimental cognitive offloading.

A study with around a thousand secondary students shows this happening in a mathematics classroom. One group practised with an unrestricted chatbot, another without. While they had it, the chatbot group scored **48 % higher on the practice problems.** Then came an exam with the chatbot taken away. Now, the same students who solved practice problems so well scored **17% worse** than the group that had never used the AI chatbot. A third group used an AI version built to withhold answers and ask questions instead. That group lost nothing.

That third group is the whole design question. A satnav that names the destination and helps you find the route teaches you the city. A satnav that says *turn left in 200 metres* does not, however good its data. Researchers call the difference guardrails, and the **design decides whether you help your students learn the city's map or simply find the destination.**

Please remember: **You have the knowledge. Your students, in the subject you teach, do not yet.** That asymmetry is the whole chapter.

Voices of the field

“Right now the tools students reach for are built to solve problems, not to teach, and this is the core difficulty in my work currently.”

Two floors, not two topics

In academic literature, these two areas are considered distinct fields. You could think of them as two floors of a building, with most teachers located on the ground floor and no staircase leading to the first floor.

In a study with 90 teachers at Bavarian vocational schools who had signed up for AI training, teachers rated their knowledge of AI at 3.4 out of 5 on average, but their ability to use it in their teaching much lower, at 2.6. Knowing and doing had come apart. A trial across five European countries then tested whether training could close that gap. 736 teachers in France, Ireland, Italy, Luxembourg and Slovenia took part. The teachers were split up in two groups. One group received an online course and support during the trial, the others received the same online course after completing the trial. The results showed that knowledge and the ability to judge what a tool could do, increased. But in most countries the regular classroom use of AI did not change: they tried tools out and went back to what they had been doing before.

This is the knowledge–action gap. Courses produce knowledge, but changing practice needs accompaniment, trialing and colleagues to think with. **Transfer has to be designed, not hoped for.**

Voices of the field
“The only drawback of this webinar is that we cannot try it out directly.”

What teachers are actually doing

The most robust finding across countries is quietly surprising: **teachers use AI intensively for themselves and hardly at all with their students**. In Germany, 52 % of teachers use AI tools several times a month or more. But of those who use AI, only 29 % let their pupils work with it in class. So the public debate about "AI in the classroom" is largely a debate about something that is not yet happening at scale.

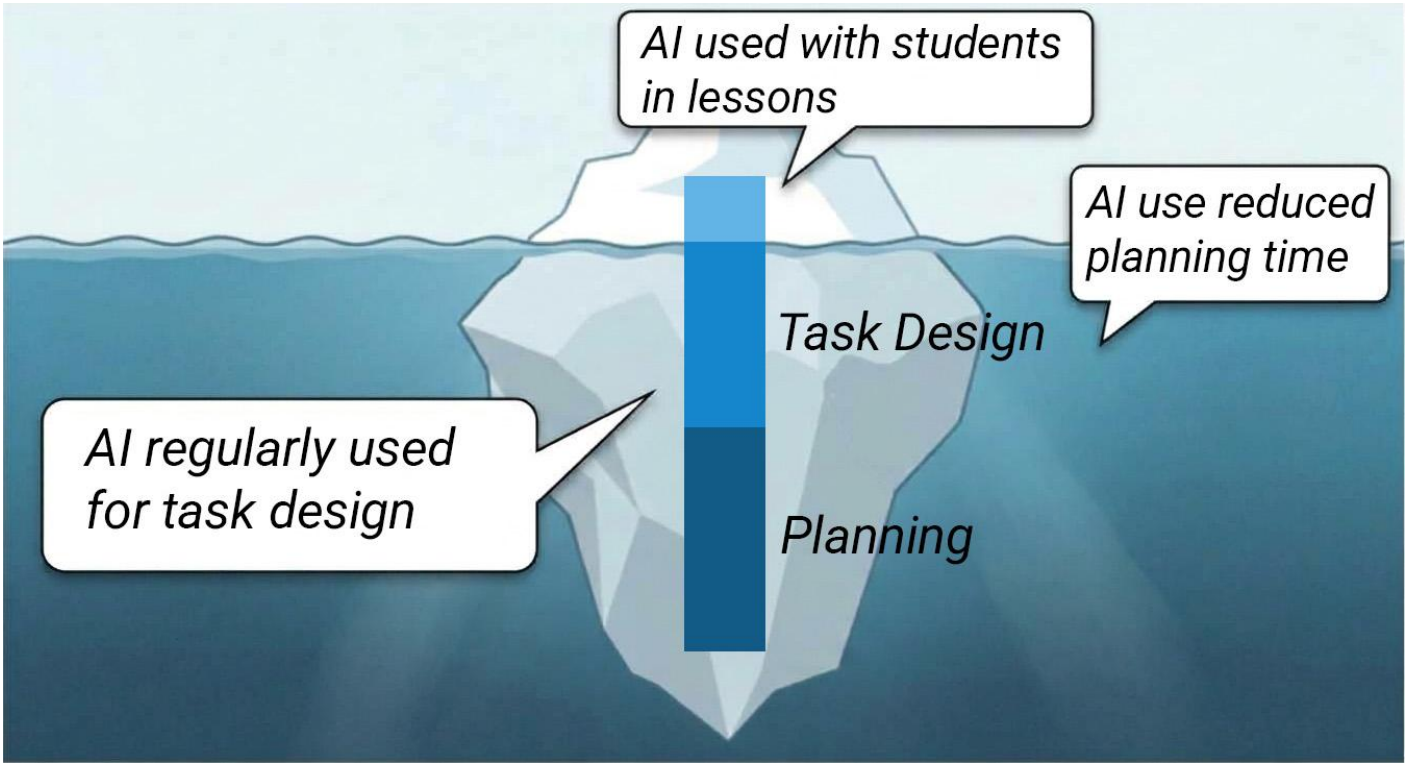


Fig. 2: Difference in AI use for preparation compared to usage in the classroom (Image generated with Google Gemini).

What **is** happening is workload relief, and here the news is good. In a trial in England, 259 teachers across 68 schools were given ChatGPT and a guide for planning science lessons for 11- to 13-year-olds. They spent **about 56 minutes a week on planning, against 81** in the comparison group. Independent assessors who did not know which group a lesson came from found no evidence of a difference in quality. In a large US survey, teachers who use AI at least weekly estimated saving almost six hours a week – their own estimate, not a measurement.

And one pattern is worth pausing on. In the same US survey, **schools with an AI policy had more teachers using AI – seven in ten against six in ten – and so saved about a quarter more teacher time in total.** For the individual teacher who uses AI every week, the difference was small. That is the teachers' own account rather than a measurement, and a survey cannot show which way the arrow runs: well-organised schools may simply do both things well. But it is the opposite of what most people expect – **clear rules and heavy use appear together, not in opposition.** **1.4 Compliant Use of AI**

Voices of the field

“It's important to have it as a supplement, but definitely not let ChatGPT run the whole lesson.”

Why the tool is not the point

If you wanted a ranking of the best AI tools, the research will disappoint you and then, on second reading, set you free. When researchers gather many separate studies and ask what separates the classrooms where AI helped a great deal from those where it barely helped at all, the answer is never the brand of software. It is what the students were asked to *do* with it. The following bullet points summarise important research results. But all results build on only a few studies.

- In mathematics, students who used AI to transform to rework a solution or to build on their own attempt gained about two and a half times as much as students who used it only to enhance ordinary tasks such as practising or checking answers.
- Students who used AI together in groups gained about 70 % more than students who used it alone..
- Where AI was used to give feedback on student work, gains were among the largest. Using AI to gamify lessons produced a gain too small to call proven.

Pedagogy is more important than the technology used. Where researchers have sorted the studies by what students were actually doing, the activity mattered more than the tool. Interestingly, the same pattern was found for the previous generation of digital tools in secondary schools. The results largely depend on whether teachers had been trained to use them.

Good feedback is not the same as feedback that works. In one experiment, 70 university students wrote an argumentative essay and revised it a week later. They were assigned at random to three groups: one got feedback from an experienced teacher, one from a chatbot given a plain instruction, and one from the same chatbot prompted to reason through the essay step by step. Marked against the same rubric, the step-by-step chatbot produced the **best feedback of the three.** The feedback was better than the plain chatbot and better than the teacher. Then came the revisions, and **all three groups improved by much the same amount.** The best feedback in the study brought no more improvement than the lowest-rated feedback. **Better comments did not produce a better essay.**

This result challenges a common assumption that better feedback leads to better learning. Obviously it does not on its own. Feedback only does anything if it is **taken up**: read, understood, acted on, and understood well enough that the student could get it right unaided next time. **Uptake is the bottleneck, not quality**.

More isn't always better. In a second study, with 60 Year 11 students and 240 pieces of writing, the AI produced way more content-related comments than the teachers did. But the students acted more often on the teachers' comments. When confronted with thirty precise corrections, a student might either stop diagnosing and start triaging or just comply. If a student follows the rules, they're basically saying that the revision was unnecessary. It's the satnav again, but on a different level. It gives you really detailed feedback, like 'Turn left in 200 metres' for your own writing.

It matters who says it. We're not really sure why teacher comments are taken up more often. This is all about interpretation, not measurement, but the most likely explanation is that feedback is seen as an act between people. What someone who knows what you did last term and will see what you do with it has to say is different from the same sentence generated on request.

What follows for practice. A pooled analysis of studies on automated writing feedback that was done before today's chatbots existed came to similar results. Feedback did more when students had to apply it to a new piece of writing than when they only revised the text in front of them, and more when it was used over several sessions rather than once. And one more thing should be said fairly: compared with no feedback at all, AI feedback does help. In a study by a German team with 459 upper-secondary pupils, AI-generated feedback modestly improved their revisions and their motivation.

And the same split shows up at scale. One team pooled 228 studies: AI clearly and consistently improved what students could produce, while any gain in their ability to plan, monitor and correct their own learning was small enough to have been chance. Programming shows it sharpest. Students with an AI assistant were measurably faster, and no better at the end than those without one. **Getting more done and getting better are not the same thing, and AI separates them.**

Europe as mirror

Across Europe, surveys of teachers' AI use run from **about nine in ten in a Czech sample to about one in seven in France.** The surveys differ in method and timing: the Czech figure comes from a small 2025 survey of 317 teachers and counts any experience with AI whereas the French figure comes from the OECD's international teacher survey TALIS 2024, in which just under half of Czech teachers reported using AI in their teaching.

France built one of the continent's largest state AI programmes: adaptive practice software in maths and French, free for every pupil in the first year of upper secondary school. Nevertheless, French teachers use AI rarely and only 9 % of French teachers have received any training in AI, compared with 38 % on average across OECD countries.

Norway has written the satnav problem into its policy (at least as a recommendation) the way a taxi driving school would do. Learn with a map first, and afterwards add the device. In June 2026 Nowrway recommended that pupils in grades 1 to 7 should, as a rule, not be given access to AI. In grades 8 to 10 AI may be introduced gradually and carefully, once teachers have the competence to guide it. In upper secondary school pupils should learn to use AI purposefully.

The blind spot

One risk of AI is rarely named, and the image reaches it too: knowledge about the London city map **is not permanent.** A driver who follows the arrows for five years does not keep the city in her head. It fades, and she stops noticing the dug-up street. Your expertise is maintained by doing the very things AI is best at absorbing: writing the task, anticipating the misconception and drafting the explanation. **Where AI takes over the work through which judgement forms, relief turns into deskilling.** In a German panel of e-learning experts, 65 out of 100 expected generative AI to cost people important skills, up from 44 the year before. In medicine the effect has already been measured: after AI support became routine in colonoscopy, doctors working without it detected fewer precancerous growths than before.

Use the satnav. Just keep the map.

1.3 Responsible Use of AI

Take Home message

- **Responsibility is a design task, not a control task.** Detection of AI does not work properly.
- **And AI detection is unfair when it fails.** Detectors wrongly flag human work as AI-written, and miss AI-written work. Education authorities in France and the Netherlands advise against using them.
- **Confident invention and inherited bias follow from how these systems are built and judged.** The consequence should be: teach them rather than waiting for a fix.
- **The working principle is: think first, then AI.** Making people commit to their own answer beats explaining the machine and building limits into the software removes the harm that unrestricted access does.
- **Prompts alone don't hold.** A chatbot tasked with acting as a tutor tends to drift back to giving answers, which makes reliable limits a purchase question.
- **Amplification is the default; equalisation must be built** from access, design and scaffolding together. The scarcest but most important of the three is a trained teacher.
- **Inclusion is the clearest win**, and where teachers themselves see the strongest case.

> If you remember only one sentence from this page: **The dams are breaking. Teach swimming.**

The reading journal, second look

Back to our ninth-grade student and his suspiciously polished reading journal. The instinct is almost universal, and it is the instinct of a good teacher: *find out whether he used AI.* Build a better dam.

This chapter is about why that instinct, on its own, leads nowhere. We will come to what works instead.

Start with the uncomfortable part. **Teachers cannot reliably tell AI text from student text.** In two studies by a German team, 89 trainee teachers and 200 experienced teachers were given a mix of student essays and chatbot essays and asked to sort them. Neither group could do so reliably but both groups were confident they were able to. Experience helped only a little: the experienced teachers still missed most of the AI texts. At one British university, researchers submitted chatbot-generated answers to real unsupervised online exams through the normal channels, without telling the markers.

Voices from the field

“You may use it, but you must be able to explain the result in your own words.”

The detector that fails the wrong students

Researchers tested seven of the most widely used AI detectors on TOEFL essays written by people learning English as a foreign language, collected from a Chinese online forum. Alongside, they put essays by native-speaking American eighth-graders. The native speakers' work was classified correctly almost every time. Of the foreign-language essays, **more than six in ten were wrongly flagged as machine-written**.

The mechanism is not mysterious. Many of these tools do not detect the work of large language models. The tools measure *predictability*. Text that uses common words in expected orders scores as artificial. Writers in a second language, with a smaller vocabulary and more rigid sentence structures, can produce exactly that signal. Meanwhile, detectors trained for one specific test, and studies in other languages, show smaller differences. Up to now there are no detectors on the market that are reliable enough to accuse anyone of willful deception.

The tool's mistakes fall systematically on the students who are already carrying the most. Institutions have done the arithmetic. Vanderbilt University pointed out that even if such a tool were wrong only 1 % of the time, that would still mean roughly **750 papers a year wrongly flagged** out of its 75,000 submissions. Consequently, they switched it off. A guide written by researchers for Germany's federal education ministry is blunt: the technology is "*not sufficiently reliable to provide legally secure proof.*" An independent European test of fourteen detectors reached the same verdict: "*neither accurate nor reliable.*"

Much of Europe has moved on from this debate. France's education ministry and its data protection authority advise schools not to use detectors, the ICT agency Kennisnet in Denmark calls them highly unreliable. In higher education individual universities, such as the ZHAW in Zurich, forbid using them for grading. In England, the exam boards allow detection software only as one element of a judgement that rests on the teacher, who "*knows their students best*". A data-protection point is worth carrying home as well: **pasting a student's work into an online tool means processing their personal data.** For this purpose, a legal basis and a contract with the provider is required. → 1.5 Compliant Use of AI

It seems that the cheating panic is thinner than the noise around it. Researchers surveyed students at three US high schools before ChatGPT existed and again afterwards. Cheating stayed broadly where it was. A follow-up survey of more than 4,000 students at six other US high schools a year later found the same.

The real question was never, "How *do I catch him*"*?* It is: **what does a homework task still measure, once its product can be generated on demand?**

Voices from the field

“At first I saw this as cheating. Now I see it as a signal that the task itself has to change.”

Two things that will not simply be fixed, so they must be taught

Confident invention is built in. These systems generate plausible continuations of a text, and plausible is not the same as true. Reviews of the research agree that no technique has yet removed the problem, although it can be reduced and partly detected. The reason is, that the standard tests used to judge large language models reward a confident guess over admitting uncertainty, so as long as they are measured that way, a certain rate of fluent, well-formed falsehood will persist.

The slants in the training material are built in too. Moreover, they do not shrink as the models improve. A European study of gender stereotyping across languages found that bigger models, even after the extra training intended to make them safer and fairer, sometimes stereotype *more* rather than less. If a property cannot be engineered away, it becomes curriculum. And it can be taught: in one Finnish programme, the share of students who could explain where an image generator's bias came from rose from **7 % to 44 %**. The consequence for teaching must be training error analysis instead of error avoidance. In other words: the machine's failures become the lesson.

Voices from the field

“Let the learner check the chatbot's solution as if they were the teacher. Looking for the mistakes is how they finally learn the method.”

What works: design, not control

The evidence points to one simple principle that's easy to remember: **think first, then AI.**

- **Make the human commit first.** In an experiment with about 200 adults, people who had to make their own decision *before* seeing the machine's suggestion went along with its wrong suggestions less often than people who were shown a short explanation of the machine's reasoning. Explaining the AI did not solve the problem. The method that helped was forcing the judgement. Forcing judgement means forcing people to think. Since thinking is exhausting the participants liked these methods least, and they helped most those who already enjoy effortful thinking. That is what your all know as the Matthew

effect: More is give to those who already have.

- **Build the limits in.** The chatbot from Chapter 1.2 , that withheld answers and asked questions instead is the same idea in software: the students who used it lost nothing on the exam, while those with the unrestricted version lost badly.
- **Make the process visible instead of hunting the product.** Drafts, working notes, a short conversation about the work. These show you what a student can do without anyone having to prove what they did.
- **Denmark is piloting the principle:** at participating schools, students may use AI while preparing for the oral English examination, and the written examination includes a handwritten, aid-free part.

Two limits are worth knowing before you rely on this.

The first: the popular advice that teachers should simply instruct a chatbot to act as a Socratic tutor may not survive a long conversation. Early evidence, from a study with simulated students that has not yet been peer-reviewed (October 2026), suggests that under persistent pressure such tutors drift back into supplying answers. Reliable limits have to be built into the system, which makes them a purchasing decision rather than a prompting skill.

The second: making an AI explain its reasoning sounds like the obvious safeguard, but explanations do not reliably prevent over-trust. In one study they left people's acceptance of *wrong* advice unchanged, and explanations that are too technical or too simplified can increase misplaced confidence. Explanation is only a condition for learning.

Attitude: two failure modes, not one

We worry that teachers are placing too much trust in AI. However, we should be equally concerned about the reverse. Even experienced people can be caught out by over-trusting. In a small pilot in US schools, four in ten of the teachers surveyed judged an AI grading tool's feedback not useful, mostly vague, or even incorrect. Research on automation bias shows that people tend to go along with a machine's suggestion even when they have reason to doubt it. **Spotting an error and correcting it are different steps.**

The opposite pole is concealment. Teachers hide their own use of AI from colleagues and students, forbidding them to do what they do themselves: 77% use AI privately, while 57% say students should never use it to generate ideas. Meanwhile, a study of academically able students in grades 6 to 8 found that between a third and two-fifths already said the AI knows more than their teachers. **A profession that hides its practices cannot model responsible practice**, and transparency is not just a courtesy. It is the mechanism by which teaching by example works.

Does AI level the field or tilt it?

For work-related tasks, AI evens out differences. In one experiment, participants with higher levels of education performed significantly better than those with lower levels of education. With an AI assistant, about three-quarters of that gap disappeared. When the assistant was taken away again, a sizeable part of the gap came back. **What the AI closed was the gap in output, not the gap in ability.**

In learning tasks, it does the opposite. One study found no average benefit at all, but a widening distance between students who already knew a lot and students who did not. Another looked at students writing in a foreign language: the weaker writers received *more* corrections from the AI and successfully acted on *fewer* of them, and several found the flood of feedback dispiriting. **Feedback without the ability to sort it does not help; it buries.**

The gap also exists before anyone touches a tool. In Germany, **80 % of students from the most advantaged homes see AI as an opportunity, against 55 % from the least advantaged**. In Britain, the divide runs through the staffroom rather than the device cupboard: **45 % of teachers at private schools have had formal training in using AI, against 21 % at state schools.** Within the state sector, wealthier schools again outpace poorer ones. A 2026 survey of more than 25,000 children across 17 European countries finds the same social split in how much and how variedly they use AI [PRÜFEN: SES-Befund am Bericht]. One commentator's line captures the whole problem better than any figure: "*the rich have access to both technology and people to help them use it, while the poor have access to technology only.*"

So the bottleneck is which learner has a teacher who has been trained. **AI compensates when three things hold together: secure access, a tool designed for learning rather than answering, and a teacher scaffolding its use.** Remove any one, and it amplifies instead. Since all three tend to be present in the same schools and absent in the same schools, **amplification is the default and equalisation is something you have to build.**

1.4 Compliant Use of AI

Take Home message

- ***There is no legal vacuum.*** Data protection law has always applied.
- ***The timetable moved.*** The high-risk duties for education AI start on **2 December 2027**, not August 2026. Check your school's handouts (see also 1.6).
- ***Already binding since February 2025.*** Systems that read pupils' emotions from their face or voice are prohibited, and schools that use AI must take measures to build their staff's AI literacy. From August 2026, chatbots must say that they are AI and deepfakes must be labelled.
- ***The hardest layer is data protection.*** No personal data about pupils without a contract that excludes training on your inputs; consent usually does not carry; deleting names is not anonymising.
- ***In copyright, the upload is the problem.*** The teaching exception does not cover feeding protected works into an AI.
- ***Detectors do not hold up as evidence.*** Courts have accepted conclusions drawn from concrete facts. Rules must be clear in advance, and asking for disclosure beats banning.
- ***Grades stay human.*** Obious to us but important to note is that rubber-stamping an AI's suggestion is not a human decision.
- ***Responsibility sits with procurement, not with you*** — *and* ***compliance enables.*** The schools with clear rules are the ones whose teachers report saving the most time.

> If you remember only one sentence from this page: **Compliance is a house, not a wall.**

What am I allowed to do?

Keep these four steps in mind before typing, pasting or uploading anything. Please use the following overview as a general guide. Requirements in your country may differ from those listed here. **Therefore, always check the regulations specific to your country and school.**

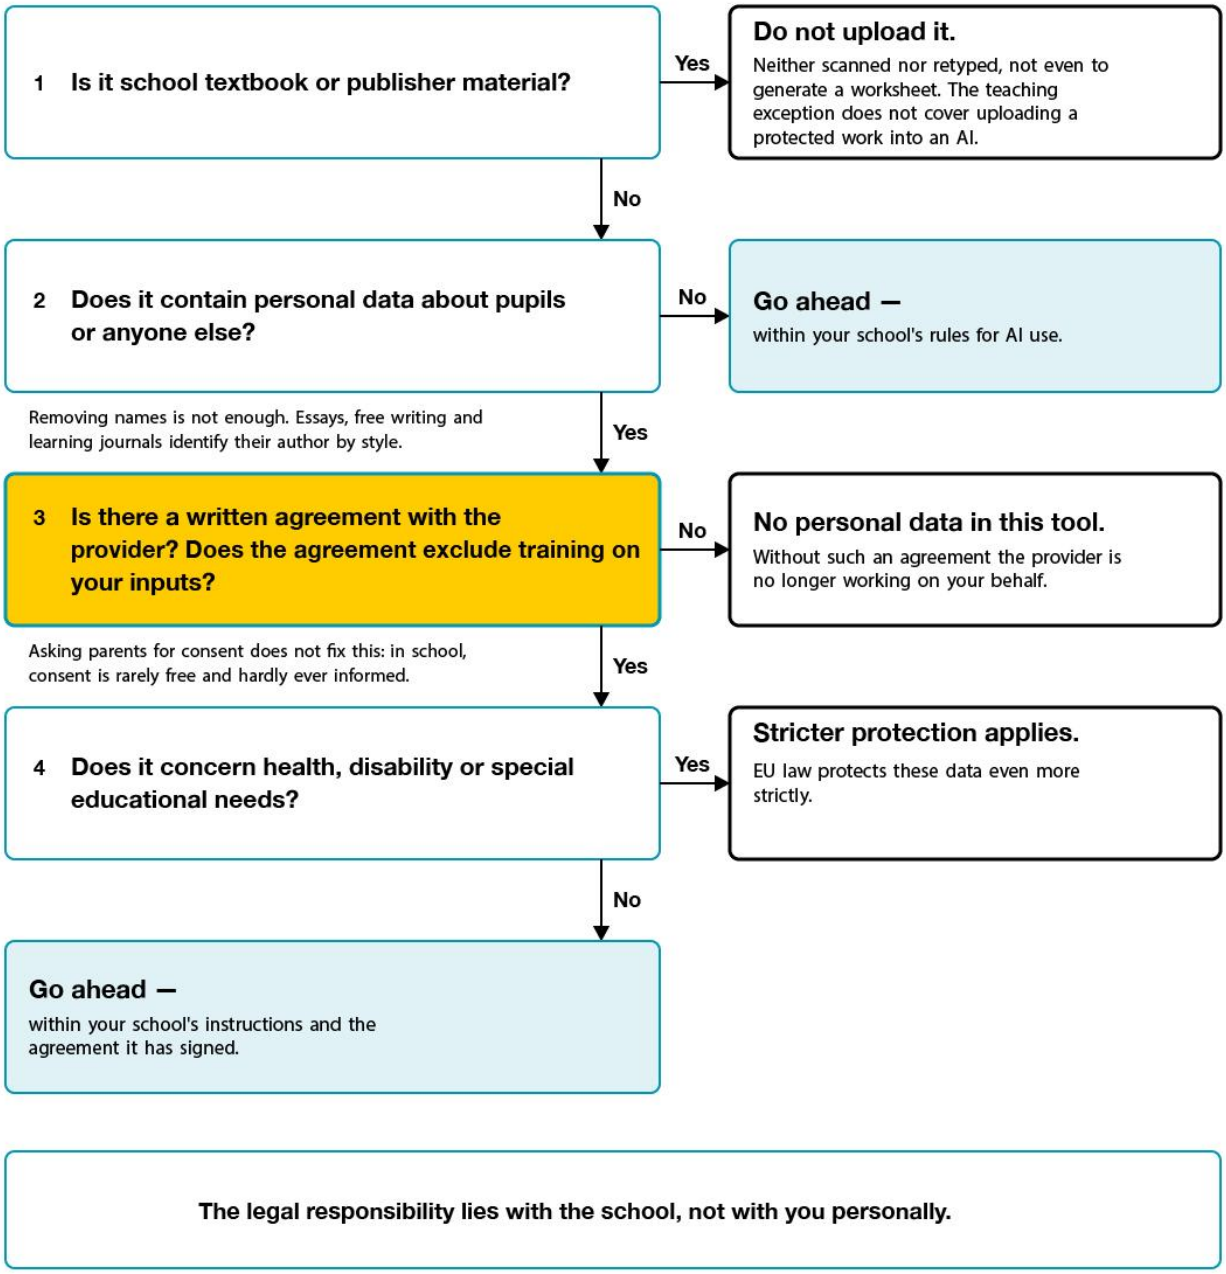


Fig. 1: ... (Image generated with ...)

Voices from the field

“We are not even sure what we are allowed to do, least of all when it comes to data protection.”

The reading journal, third look

Our ninth-grade student, his reading journal, his teacher's suspicion. In Chapter 1.2, it was a question about learning; in Chapter 1.3, about professional judgement. Here it is a court file.

In December 2025 the Administrative Court of Hamburg, Germany, ruled on an urgent application in the case. The English teacher had noticed that the reading journal was written at a far higher level of spelling, grammar and vocabulary than the boy's class test on the same book. He later admitted using ChatGPT, and the school graded the journal as an attempt to deceive. His father applied to the court for urgent relief; the court refused. Its reasoning goes beyond this case: pupils must assume that they are to do their work entirely on their own unless an aid has been expressly allowed – so undisclosed AI use counts as deception even without a written rule against AI. (In this case, the teacher had in fact told the class that ChatGPT was not allowed.) One caveat to keep in mind from the outset: this is a single court, in urgent proceedings, and the decision is not final.

This is a good moment to address the most common misconception in this field. **There is no legal vacuum.** There never has been.

The house has three floors

Most teachers envisage AI law as a single wall under construction somewhere in Brussels. However, it is more useful to picture a house that you already live in. **The foundation is data protection.** It is load-bearing. It was laid years ago. It governs your Tuesday morning. **The upper floor is the EU AI Act,** Europe's AI law, which categorises systems by their potential for harm and imposes stricter obligations as risk rises. It is partly occupied and still under construction, with the completion date just pushed back. **The doors and windows are copyright and examination law** — small and specific, and touched daily without noticing.

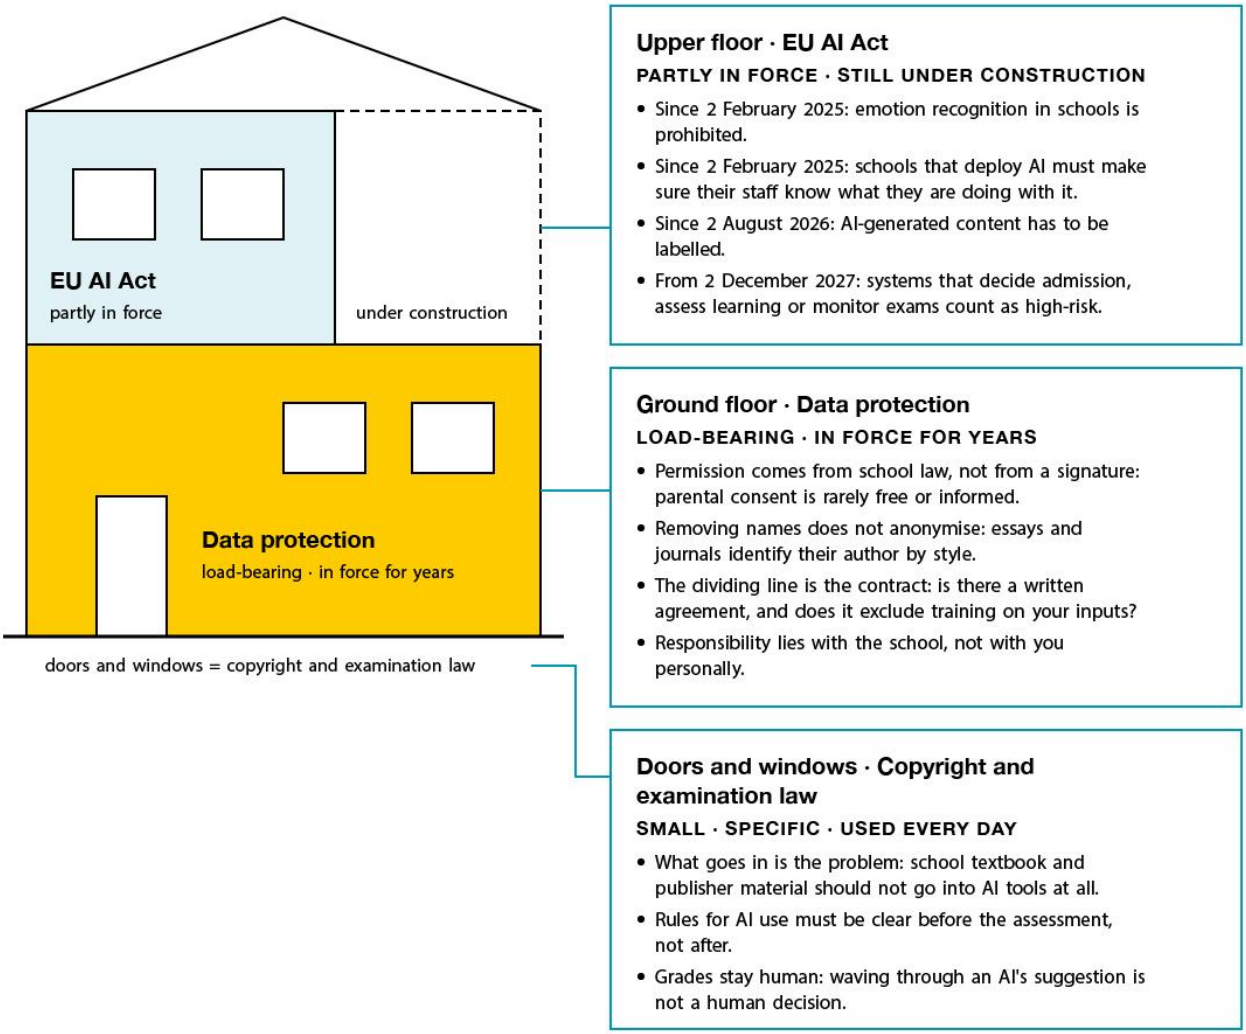


Fig. 2: Compliance is a house (Image generated with ...)

Ground floor: data protection, clearer than many think

Handling someone's personal data always requires legal permission. As a public body, a school's permission essentially comes from its statutory duties under school law. Being a teacher does not come with a general licence.

Three consequences follow, forming the practical core of this chapter.

Consent does not provide a solution. One obvious solution would be to simply ask the parents. However, this approach has two drawbacks. Firstly, consent must be freely given. In a school context where grades are a factor, a teacher's request does not constitute a free choice. Second, consent must be informed, which is difficult when nobody knows where the data will end up. **Permission must come from school law**, not a signature, as this shifts responsibility upwards, away from you.

Removing names does not anonymise data. In 2024, Germany's data protection authorities jointly stated that if someone could still be identified from the context, the data is still personal. This is particularly problematic in schools because the types of text that teachers most want help with — such as essays, free writing and learning journals — identify their author by style.

The dividing line is the contract: a provider may handle your data if it does so strictly on your instructions and there is a written agreement to this effect. However, the moment it uses what you type for its own purposes, the arrangement collapses because it is no longer working on your behalf. So the practical test is this: I**s there such an agreement, and does it exclude training on your inputs?** If not, no personal data should be entered.

Two footnotes worth noting: data relating to health and disability is subject to even stricter protection under EU law. It does not belong in an AI tool. And where teachers use private devices, at least two German states (North Rhine-Westphalia and Bavaria) make clear that **the data-protection responsibility lies with the school and its head**, who must approve the use.

Upper floor: what the AI Act does and does not require yet

Here is the correction that makes most existing guidance out of date. In July 2026, with an amending regulation known as the Digital Omnibus, the EU postponed the heaviest set of AI Act duties for education which are those attaching to systems classed as **high-risk**, meaning systems used to decide admissions, to assess what pupils have learned, or to monitor them during examinations, **from August 2026 to 2 December 2027**. A great many handouts written in 2025 and early 2026 still print the old date. If a document in your school does, it needs a footnote.

The more useful correction runs the other way: **what already applies is underestimated, and what is endlessly discussed does not apply yet.**

Two things have been binding since **2 February 2025**.

- **Emotion recognition in schools is prohibited.** Systems that claim to read pupils' feelings from their face, voice or body – mood analytics, "engagement" scoring from a webcam – sit on the Act's short list of practices that are simply banned, with narrow exceptions for medical or safety reasons. Tools that only detect tiredness or whether someone is looking at the screen are a grey area; that is a question for procurement, not for the classroom.
- **Schools that deploy AI must take measures to build their staff's AI literacy.** Since July 2026 the law states that no particular level has to be guaranteed – but doing nothing is not an option. There is no EU-level fine for this duty, which is why it is widely overlooked; national authorities can enforce it from August 2026.
- **From 2 December 2026, AI that creates sexualised deepfakes of identifiable people is prohibited.**

Read that second duty alongside Chapter 1.3 and something clicks. The gap between knowing about AI and being able to teach with it is no longer only a professional problem. **Since February 2025, schools have been legally required to work on it.**

Since **2 August 2026**, chatbots must tell users that they are AI, and providers must mark AI-generated content in a machine-readable way; schools themselves must label deepfakes, and AI-written texts they publish to inform the public. A worksheet a teacher has checked needs no label. From **2 December 2027**, the duties for high-risk systems apply – systems used to decide who gets admitted, to assess what pupils have learned, or to monitor them during examinations. Whoever *uses* such a system then takes on duties of its own, including assigning oversight to people with, in the Act's words, "*the necessary competence, training and authority*".

Doors and windows: copyright and examinations

Counter-intuitively, **what comes out is the easy part.** A text or image that an AI generates on its own has no human author in the sense copyright requires, so it is usually free of copyright itself; where a person shapes it through their own creative choices, the result can be protected. It is not, however, guaranteed to be clean: in a widely noted ruling, a Munich court found that a chatbot had memorised song lyrics and reproduced them almost word for word. The court placed the liability on the **provider**, not on the person who typed the request. The judgment is not final.

The exception that lets teachers copy material for their classes covers exactly that: copying a limited share of a work, for that class. **Uploading a protected work into a commercial AI system is a different act, and the exception does not address it.** Material produced specifically for school teaching, such as textbooks and workbooks, is in principle excluded from the exception, as long as licences for it are readily available.

How much should that worry you? Less than the gap suggests, and more than nothing. The question is **genuinely unsettled**: no higher court has ruled on it, and the specialist literature regards the existing exceptions as a poor fit for the situation. Where a rights holder objects, the claim would be a civil one, and in practice, this is a matter for school policy and procurement rather than something an individual teacher can settle. What is clear in practice is the narrow case. It is the line the educational publishers draw, and it is the one worth remembering: **school textbook and publisher material should not go into AI tools at all — neither scanned nor retyped, not even to generate a worksheet.**

Case law in this field is currently inconsistent across Europe. In February 2026, the Administrative Court of Kassel, Germany, decided two university cases the same way, and sharply: even though the university had no explicit ban on AI, the line was crossed "*already with a single undisclosed use of generative AI*". The judgments are not final. In the same month, a Paris administrative court, in urgent proceedings, stopped a university from sanctioning a student because it had set out no rule on AI use in advance [PRÜFEN]. A Swedish appeal court had reached a similar result earlier: a student had handed in three invented quotations, but the course rules allowed "all aids" in take-home work, so there was no rule to have broken. In the Netherlands, by contrast, the country's highest administrative court upheld a sanction in July 2025 against a law student who had submitted four non-existent sources: the fraud was proven beyond reasonable doubt, and whether he had meant to deceive did not matter.

So the pattern across jurisdictions is not about countries. It is about **whether a rule existed beforehand, and whether the deception was actually shown.**

Three rules follow for practice.

1. **Detectors do not carry the burden of proof.** The school has to show the deception. Courts have accepted conclusions drawn from concrete facts: sources that turn out not to exist, a clear gap between what a student writes and what they can explain orally, sudden jumps in level within the same text, an admission. The student can then explain. Writing style alone is a weak basis – leading German legal commentators doubt that it can carry a finding at all. Detector output is, at best, a reason to look more closely. The failure rate reported in Chapter 1.3 makes it unusable as evidence.
2. **Rules must be clear before the assessment, not after.** Neither a blanket ban nor a blanket permission holds up. What works is a declaration of independent work with a clause about AI, plus a duty to state what was used. Disclosure instead of banning.
3. **Grades stay human.** People have the right not to be subject to decisions with legal effect that a machine takes alone, and schools have no legal basis for making such decisions by software; a grade is one. Germany's Standing Conference of the Ministers of Education and Cultural Affairs puts it even more strongly: assessment is a task only teachers can perform. Crucially, **formally waving through an AI's suggestion does not count as a human decision.** So the finding that people often go along with flawed AI feedback is a legal problem as well as a pedagogical one.

Voices from the field
“Everyone is experimenting on their own, so the rules change from one lesson to the next.”

Law, ethics and didactics therefore converge on the same answer: **assess the process, not just the product.**

The real address of responsibility

In response, Germany has introduced state-provided school AI environments. These are vetted platforms with negotiated contracts, rather than leaving teachers to sign up to whatever is free. While these solve problems that individual teachers cannot solve alone, they are still criticised for how they are governed, whether teachers actually accept them and whether the assessments they support are sound. Because the rollout differs from one federal state to the next, it will take time to see how teachers are faring with the platforms.

What is distinctive is narrower than it may initially appear, and it is worth stating precisely. Germany is the only system we found in which a public body – the FWU, owned by the federal states – develops and runs the platform itself, with technical contractors. France provides adaptive practice software for an entire year group, but commissioned it: one private company wrote the software, another hosts the data. Several other countries have bought access to commercial chatbots; most of the rest offer guidance only.

Building it yourself does not automatically give you more control. The German platform runs on the same commercial language models as everyone else's on European servers, as an industry association has pointed out. By contrast, the French system is hosted in France under the highest security certification issued by the French state. One country built the surface and rents the engine, while the other bought the software and keeps the keys to the building. Neither option is clearly the most sovereign choice, which is worth bearing in mind before using the term.

The Netherlands took a different approach, and it is worth knowing about because it offers a more honest answer to the same question. Rather than building or buying, Dutch schools and universities joined forces and had the major providers formally assessed. When they examined Microsoft's AI assistant, they found two issues that could not be resolved: the criteria for an automatic content filter that the provider would not disclose and the retention of eighteen months of diagnostic data that could not be justified. They negotiated. In March 2026, however, the provider declined to change. The assessment therefore remains amber, and Dutch schools are advised to exercise caution and use their own judgement on a case-by-case basis rather than granting clearance. One country removed the risk by building around it, while another took a closer look and decided it was not ready yet. Both approaches offer solutions to the problem set out at the start of this chapter.

The comparison also settles one more thing. The complaint that teachers prefer the free chatbot to the official one is not due to a failure of implementation in Germany. In Estonia, where the state negotiated access for upper-secondary students, barely more than a third were using the official chatbot weekly seven months after the programme began. It turned out that most of them had been using free tools before it started. **Everywhere, the vetted environment competes with the more convenient one.** Providing the official tool is necessary. But this is not sufficient.

However, the structural point stands: compliance is an enabling condition, not a hindrance. Remember what was said in Chapter 1.2? **Schools with clear rules had more teachers using AI – and saved more time overall.** Rules did not slow use down

1.5 Current State of AI in Education

Take Home message

- **AI is already widely used in educational contexts.**
  - **Teaching with AI.** AI is currently used by teachers to prepare or improve lesson plans, for assessment and to distinguish content across different educational levels.
  - **Teaching about AI.** AI is used by students though they have no groundings on what is AI, which are its limitations and to use it responsibly.
- **Technical understanding and critical thinking are needed.** Teachers and students need to understand how language models work (training, data handling, predictive nature) and develop practical verification skills, from spotting fabricated citations to fact-checking AI-generated content.
- **Responsible use and ethical stance are to be defined**. A grasp of the legal and ethical landscape (AI Act, GDPR) is needed, along with an approach to AI as a support rather than a substitute for human thinking, grounded in transparency and accountability, plus explicit instruction in prompting as a teachable skill.
- **Equity and new assessment models are needed.** Equitable access to AI tools needs safeguarding (public funding, open models, minimum age policies), alongside a shift in assessment from grading the final product to evaluating the process and students' ability to critically engage with AI output.

> If you remember only one sentence from this page: **AI is already present in educational contexts but is not regulated. Structured frameworks are needed to support educators' professional development improving their knowledge and letting them teach with AI and about AI.**

Current Uses of AI in Educational Contexts

Nowadays, artificial intelligence (AI) is embedded in everyday educational practice, well beyond pilot projects. Common applications include adaptive learning platforms, intelligent tutoring systems, automated grading, and AI-assisted lesson planning for teachers (OECD, 2026). The most visible shift, however, is the direct use of general-purpose tools by students and teachers alike.

Teaching With AI

AI is increasingly used by teachers as a support tool across multiple stages of their work, from lesson planning to grading. According to OECD's TALIS survey, 37% of lower secondary teachers reported using AI in 2024, with 57% saying it helps them write or improve lesson plans (OECD, 2026). Beyond lesson design, teachers are using AI to differentiate content for different skill levels, generate practice materials, and support the grading process. At the same time, adoption is not without friction: 72% of teachers expressed concern that AI could enable students to pass off AI-generated work as their own, reflecting a broader tension between AI's practical benefits for teaching and its implications for academic integrity (OECD, 2026).

Teaching About AI

Alongside using AI as a teaching aid, there is a growing need to teach students about AI itself: how it works, its limitations, and how to use it responsibly. This includes building a basic understanding of language models as prediction engines rather than databases, which explains their tendency to hallucinate, as well as practical skills such as fact-checking AI output and recognising fabricated citations (UNESCO, 2024). It also involves developing critical thinking around prompting, shifting students from simply following instructions to formulating them, and instilling awareness of the ethical and legal dimensions of AI use, including data privacy and the EU AI Act. UNESCO's *AI Competency Framework for Teachers* (2024) frames this as a core professional competency, arguing that teachers themselves need training in these areas before they can effectively teach them to students.

What is missing?

Alongside institutional frameworks and adoption data, it is important to understand how these dynamics play out in practice. A series of workshops was conducted in the context of the DUAL.AI.TEACHer project. Numerous education stakeholders were called to gather first-hand perspectives on the current use of AI in teaching. These workshops surfaced a set of concrete needs and concerns expressed by educators, covering practical, pedagogical, and training-related aspects. The following sections outline these needs in detail.

- **Understanding How AI Works.** Both students and teachers need to understand how language models are built and trained, and what happens to their data once provided. They also need to grasp the nature of these models as prediction engines rather than databases, which explains their tendency to hallucinate. Teachers need a simple high-level course covering the basics of how LLMs work.
- **Critical Thinking, Questioning and Verifying AI.** Students and teachers need to be able to question AI-generated answers and to recognise fabricated information. Practical verification skills are also needed, from checking AI output against official sources to spotting AI-generated images and videos. Students need to shift from following instructions to formulating them, thinking critically about the desired output.

- **Ethics, Data Protection and Regulation.** A general understanding of the ethical and legal aspects of AI use is needed, including the EU AI Act and GDPR. Safe and responsible use is essential to prevent illegal practices along with awareness of related issues like bias in teaching materials and the unreliability of AI detectors.
- **Responsible Use.** AI needs to be used mindfully as a support, not a replacement, leaving room for students' own thinking. Transparency matters, with students expected to flag which parts are their own, echoing "cyborg writing", where AI extends rather than replaces the writer's thinking. Student ownership and accountability are also needed.
- **Practical Skills.** Prompting needs to be treated as a teachable skill, not something acquired independently. Both teachers and students need explicit instruction in writing effective prompts, and teachers need training in prompt engineering for lesson planning, differentiation, and grading. Broader tool familiarity is also needed to match tools to subjects.
- **Equity and Infrastructure.** Access to AI is at the centre of growing equity concerns, as wealthier districts can afford better tools, widening competence gaps between schools. Public funding is needed to guarantee equal access regardless of a school's resources, preferring open AI models over corporate ones as part of this effort. A minimum age for AI use also needs to be defined.
- **Assessment.** A different approach to assessment is needed, focused on process rather than final product, including students' reflections on AI-generated drafts. AI-integrated assignments are also needed, grading students on their ability to critique, fact-check and edit AI output, raising questions about the human role when teachers themselves rely on AI for grading. Concrete classroom methods are needed too, such as fact-checking AI text, comparing chatbot responses, debating AI use in homework, and incorporating non-digital tools like flip charts.

Voices from the field

“Anyone who teaches with AI must be able to verify every answer it gives, because errors happen. Without solid subject knowledge you will not spot them.”

1.6 Legislation

Take Home message

- **The EU AI Act is the legal basis.** The EU's rulebook for AI is in force since August 2024 and phasing in through 2027. It applies not just tech companies, but also to schools.
- **AI literacy is a duty, not a course to tick off.** Schools must support staff in building AI literacy.
- **Several AI uses are banned outright.** Among them: inferring students' emotions from facial or vocal data, AI that exploits children's vulnerabilities to distort their behaviour, and building facial-recognition databases or inferring traits like race or religion from biometric data.
- **Risk depends on the decision, not the product.** Ordinary writing, research, translation and lesson-planning tools are not high-risk simply because they are used at school. They become high-risk once they influence admission, learning pathways, exam outcomes or employment decisions.
- **Not everything AI-generated needs a label.** The disclosure duty mainly sits with the provider; teacher-reviewed material with human editorial control is exempt.
- **High-risk AI needs a human who can overrule it.** AI output is never the final decision. You must understand its limits and be able to disregard or override it.
- **The AI Act sits on top of other law, it doesn't replace it.** National data protection, education, employment, child protection and copyright law continue to apply alongside it.
- **The GDPR governs the data, not just the tool.** Even an AI system that is fine under the AI Act can still be off-limits if it processes personal data about students or staff without a valid legal basis.

> If you remember only one sentence from this page: **Use AI to support your professional judgement, not replace it—stricter rules apply as soon as AI helps make decisions about grades, educational pathways or people.**

The EU AI Act

The EU AI Act (Regulation (EU) 2024/1689) is the EU's single, risk-based rulebook for artificial intelligence. Instead of regulating specific products, it classifies AI systems by the risk they pose and attaches obligations accordingly — from an outright ban for a small set of practices, through extra duties for "high-risk" uses, to no special rules at all for everyday tools. It applies directly in every member state, to public and private organisations alike.

The Act phases in step by step. It has applied since 1 August 2024; since 2 February 2025, the banned practices below and the duty to build AI literacy (Art. 4) are already binding; from 2 August 2026, most of its remaining provisions apply.

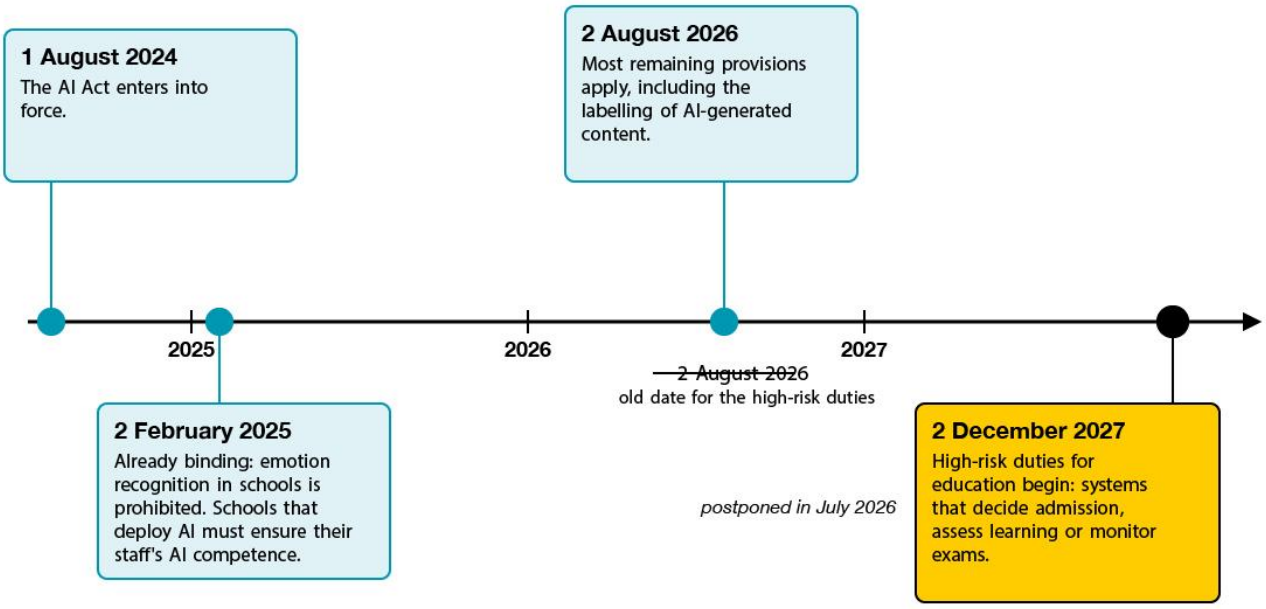


Fig. 1: The "AI roadmap" (Image generated with ...)

Who is responsible for the use of AI?

If a school introduces and controls an AI system, the school or school authority is usually the *deployer* (Art. 3(4)). The teacher is normally not a separate deployer when using the tool under the school's instructions, but must still follow those instructions and all relevant school and data-protection rules.

Teachers need to know what they are using

Schools and school authorities must take appropriate measures to support the development of AI literacy among those who use AI systems on their behalf (Art. 4). The AI Act does not specify a particular level of competence, certificate or training format. These rather depend on the AI system used, how it is used, and who may be affected.

AI practices banned outright for schools

The AI Act bans eight AI practices outright (Art. 5); four of them are directly relevant to schools:

1. Using AI systems in educational institutions to **infer emotions or intentions from biometric data** (such as students' facial or vocal characteristics) is generally forbidden. Narrow exceptions exist for medical or safety purposes (Art. 3(39); Art. 5(1)(f)).
2. AI systems that exploit someone’s vulnerabilities due to age (as with children), disability, or social or economic situation, in a way that distorts their behaviour and causes them harm, are banned regardless of intent (Art. 5(1)(b)).
3. Building or expanding a facial-recognition database through untargeted scraping of images from the internet or from CCTV footage is banned. This is relevant if a school considers a facial-recognition tool for attendance or exam monitoring (Art. 5(1)(e)).
4. Using biometric data to categorise people in order to infer protected traits such as race, political opinions, religion or sexual orientation is banned. This is relevant for any biometric-based tool, such as exam-proctoring software, that goes beyond its stated purpose (Art. 5(1)(g)).

Does AI-generated content always need a label?

If students interact directly with an AI system, such as a chatbot, they must be informed that they are interacting with AI. Usually the system provider is responsible for making this clear (Art. 50(1)).

Schools and teachers must disclose deepfakes, or AI-generated text published on matters of public interest, when these are used in the classroom (Art. 50(4)). This requirement does not apply where the AI-generated content has undergone effective human review or editorial control, and a person or institution assumes editorial responsibility for its publication. An AI-assisted worksheet reviewed by a teacher therefore does not automatically require a label.

When AI in education becomes high-risk

Ordinary writing, translation, research and lesson-planning tools are not high-risk simply because they are used for work. Classification depends on the system's intended purpose, not simply on which product is used (Art. 6 · Annex III).

AI systems may be classified as high-risk when they are intended to:

For Teachers	For Schools
• decide about students' access, admission or assignment to educational or vocational institutions and programmes; • evaluate students' learning outcomes or steer their learning processes; • assess the appropriate level of education a student will receive or can access; or • detect students' prohibited behaviour during examinations. *(Annex III, point 3)*	• recruit or select teachers, including screening applications; • decide on contract terms, promotion or termination of a teacher's employment; • allocate tasks based on a teacher's behaviour or personal traits; or • evaluate teachers' performance. *(Annex III, point 4)*

A limited checking or preparatory tool may also fall outside the high-risk category if it does not significantly influence a decision. However, an education system that profiles individuals is always considered high-risk.

What high-risk AI means for schools and teachers

Classification of an AI system as high-risk does not automatically rule it out from use in schools. But high-risk systems come with obligations (Art. 26):

These obligations for high-risk AI in education apply from **2 December 2027** — but the prohibitions above and the AI-literacy duty already apply today.

For Teachers	For Schools
• use it only for its approved purpose and follow its instructions; • do not treat its output as a final decision; • understand its main limitations and likely sources of error; • watch for overreliance on apparently convincing results; and • be prepared to disregard or override its output.	• provide competent and authorised human oversight; • ensure that any input data under its control is relevant and sufficiently representative; • monitor the system and respond to risks or serious incidents; • keep automatically generated logs for at least six months where those logs are under its control; and • inform students, teachers or affected people in advance when the system assists decision-making about them.

Public bodies and private organisations providing public services may also have to complete a fundamental rights impact assessment before using high-risk AI (Art. 27). Public authorities and organisations acting on their behalf may additionally have to register the system in the EU

database (Art. 49). Whether this is the responsibility of the school or its governing authority depends on the national school system.

The AI Act is not the only law that matters

The AI Act does not replace european or national regulations on data protection, education, employment and worker participation, child protection or copyright (Art. 2(7)). For example, the fact that a tool is not high-risk does not mean that you may upload students' work or personal data to it.

A tool may be acceptable under the AI Act but still prohibited by data-protection law or your school's rules.

The GDPR

The GDPR (Regulation (EU) 2016/679) is the EU's law on data protection. It plays a supporting role alongside the AI Act: while the AI Act asks whether and how an AI system may be used at all, the GDPR asks a narrower question — may this particular data about this particular person be processed this way? Meeting one does not automatically mean you have met the other.

- **It applies as soon as personal data is involved.** Names, school email addresses, student work, voice recordings, transcripts, photos, grades, attendance and learning-progress data all count as personal data (Art. 4). Most AI tools used in class touch at least one of these.
- **You need a lawful basis to process it.** A legal obligation or a task carried out in the public interest are the bases schools usually rely on; consent is also possible but must be freely given and, for younger students, given by a parent or guardian (Art. 6, Art. 8).
- **Some data needs extra protection.** Health information, biometric data and data revealing racial or ethnic origin, religion or similar traits are "special category data" and are, in principle, off-limits without a specific exception (Art. 9) — one reason the AI Act separately bans several biometric-inference practices in schools.
- **Know who is responsible.** The school is normally the "controller"; an AI provider processing data on the school's behalf is a "processor" and needs a data-processing agreement that spells out its obligations (Art. 28). Check this before rolling a new tool out school-wide.
- **Data leaving the EU needs a safeguard.** Many popular AI tools are hosted outside the EU. Personal data may only travel there if an adequacy decision or another approved safeguard, such as standard contractual clauses, is in place (Art. 44).
- **Students and staff keep their rights.** They can ask what is held about them, have it corrected or deleted, and object to decisions based solely on automated processing that significantly affects them (Art. 15–Art. 22) — a right that runs alongside, and reinforces,

the AI Act’s human-oversight requirements for high-risk systems.

In practice: before feeding real student or staff data into an AI tool, check that your school has approved it, share no more personal data than the task needs, and prefer anonymised or fictional examples when just trying a tool out.

Where to find help

Enforcement is organised differently in each EU country. Because national responsibilities may change, use the European Commission's current list of national market-surveillance authorities. Anyone (a teacher, a student, a parent) can lodge a complaint with that authority if they believe the Act has been breached (Art. 85). Where a high-risk system, such as an admission or exam-monitoring tool, has led to a decision that significantly affects a student, they (or their parents) can also ask for an explanation of that individual decision (Art. 86).

National rules for your country

The AI Act is EU-wide, but enforcement bodies, school law and data-protection practice differ by country. Country-specific notes for the Playbook's partner countries will be added here:

- Austria
- Czechia
- Germany
- Latvia
- Poland
- Slovakia
- Slovenia
- Spain