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.**

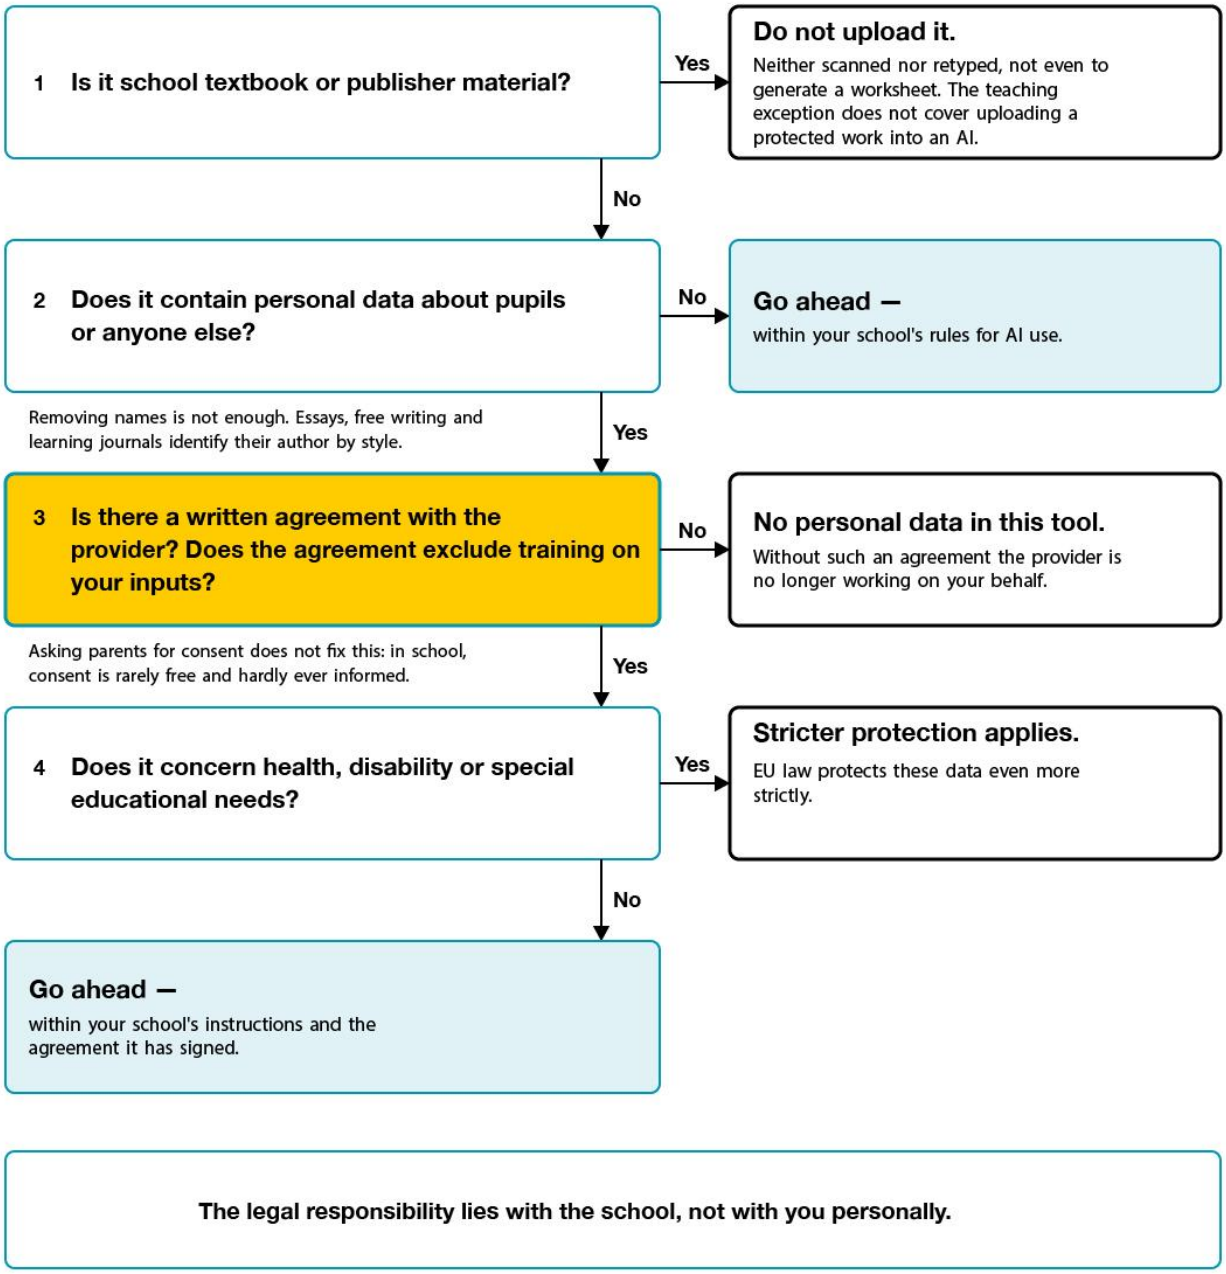


Fig. 1: ... (Image generated with ...)

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.

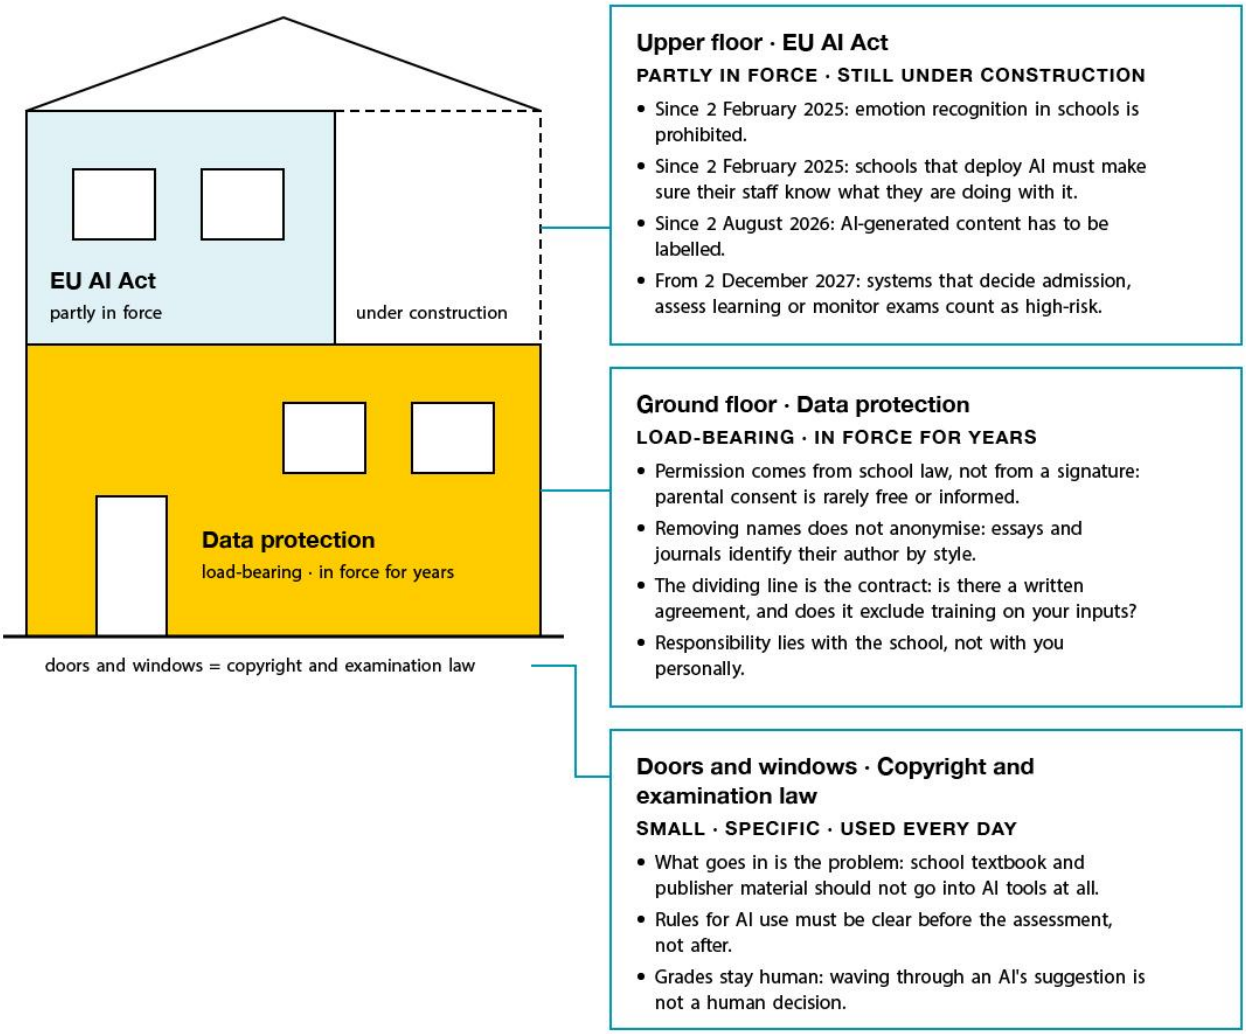


Fig. 2: Compliance is a house (Image generated with ...)

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

Revision #8
Created 2026-10-20 11:34:13 UTC by Jannik
Updated 2026-10-20 11:34:15 UTC by Jannik