ACTION PAPER · MAY 2026

AI cheating in K–9 schools

A practical guide for teachers and school leaders on assignment design, the conversation with the student, and the official framework — for Danish K–9

TEACHERS & LEADERSHIP GUIDE MAY 2026

Contents

About this document 3

1. The new reality 4

2. AI detectors do not work 5

3. When a submission could be AI-written — and why that is not proof 6

4. Seven types of assignment that work with AI 7

5. “I have 28 students — this will take forever” 10

6. The conversation when you suspect a student 11

7. The official Danish framework 12

8. Template: local policy for AI use 13

9. Do and don’t 14

10. The most important thing to remember 15

About this document

This is a short, practical guide for teachers and school leaders in Danish K–9 schools facing the same question as everyone else: *the students are using AI for their assignments — what do I do about it?*

The guide gathers the most robust findings from research and practice up to May 2026. It is not a legal handbook. It is not neutral. It commits to one position — that the problem is best solved through assignment design and dialogue, not through surveillance and detection — and builds from there.

A note on terminology: *AI cheating* refers to what the student and the school would jointly call cheating: submitting AI-written content as one's own. *Detector* covers tools such as Turnitin AI and ZeroGPT that attempt to determine whether a text was generated by AI. *Generative AI* covers ChatGPT, Claude, Gemini and similar language models. The document is centred on Denmark, where the Folkeskole covers compulsory schooling for ages 6–15 (grades 0–9). I use "K–9" as shorthand for the same age range throughout.

Who should read what?

- **School leader or subject coordinator:** Start with **section 7** (the official Danish framework) and **section 8** (template for a local policy). Then return to **section 1** (the new reality).
- **Teacher in grades 1–4:** Read **sections 1, 4** (assignment types) and **section 5** (the practical day-to-day).
- **Teacher in grades 5–6:** Read **sections 1, 4, 5** and **section 6** (the conversation with the student).
- **Teacher in grades 7–9:** Read the whole document. This is where the problem is most acute.

1. The new reality

AI is not a fad that will pass. RAND data from the US shows that the share of students using AI for homework rose from 48 per cent to 62 per cent during 2025 alone. In total, 92 per cent of students use AI for schoolwork in one form or another. The figure comes from an American panel, but the Danish pattern is closer than many imagine — figures from the Danish Evaluation Institute (EVA) from 2025 show that two thirds of upper-secondary final-year students have used AI in a way they themselves call cheating.

You therefore face a binary choice each time you design an assignment:

- **Resist.** Spend time detecting AI use, introduce surveillance, frustrate the students, lose the energy that makes teaching meaningful.
- **Adapt the assignments.** Design submissions where AI is either irrelevant, openly permitted, or impossible to hide.

This is not a parallel to the calculator, as is sometimes claimed. A calculator takes over *the calculation*, but the student must still understand the mathematics to use it. A generative AI can take over *thinking itself* — the analysis, the synthesis, the writing process. A study in *npj Science of Learning* (2025) suggests that students who use AI intensively score lower on critical thinking over time.

The conclusion is not that AI should be banned. It is that assignments should be designed so the student *still has to think for themselves* — even when AI is available.

2. AI detectors do not work

The first question many ask is: can I use a detector to expose AI use? The short answer is no. This has been examined in four studies and a substantial review since 2023, and the conclusion is consistent.

False positives are too frequent

Turnitin AI reports a false positive rate of around 3.8 per cent. ZeroGPT runs as high as 16 per cent. None of the available detectors is reliable enough to justify a false accusation — and you have no way of knowing which student will be the one in twenty-five who is wrongly flagged.

Students can easily circumvent them

The prompt *"Rewrite this text so that it is not detected by AI detectors"* works. On lightly edited or so-called *humanised* AI text, most detectors drop below 40 per cent accuracy. A 2026 review estimates that 94 per cent of AI-generated assignments now pass detection undetected.

The detectors discriminate

A Stanford study showed that AI detectors falsely flagged 61.3 per cent of English essays written by Chinese TOEFL students as AI — against only 5.1 per cent for essays by American students. Bilingual students often write "differently from normal", and detectors mistake this for AI. In a Danish classroom, this means bilingual students are disproportionately hit by false accusations.

Watermarking will not save us

Google places invisible *watermarks* on AI text from Gemini (SynthID). OpenAI has built — but never released — an equivalent system. In practice, none of the major tools marks its output reliably enough to function as evidence.

Conclusion: Detectors create false reassurance and produce uneven outcomes. Focus instead on assignment design and the conversation with the student.

3. When a submission could be AI-written — and why that is not proof

You will still often have a hunch. Some signals have historically pointed towards AI use, but it is important to recognise that they are *conversation starters*, not verdicts.

- **A sudden change in language level.** The student normally writes simply, but submits a text with noticeably higher language than before.
- **The student cannot explain their own work.** You ask the student to unpack a passage. If they struggle to explain what they themselves have written, that is an indication.
- **The content is generic.** AI often writes correctly but blandly. The personal angle, local references, and the student's own experience are missing.
- **Good grammar, weak argumentation.** AI is strong on surface language, weaker on original thought and independent argument.
- **No personal experience.** The assignment asks for examples from the student's own life, but the text stays at an abstract level.

Two important caveats:

A signal is not proof. Many students have simply got better all of a sudden. Others have had help at home. Bilingual students often write "differently from normal" without using AI. Never accuse a student without speaking to them first.

The signals weaken year by year. GPT-5, Claude 4.x and Gemini 2.x write more informally and personally than ChatGPT-3.5 did in 2023. *Format consistency* — tidy bullet points, clear headings, structured text — was a stronger flag against the 2023 models. Today, students can ask the AI to "write like a year-7 pupil", and it will do so reasonably well.

The signals are still useful as *grounds for a conversation* — not as grounds for an accusation.

4. Seven types of assignment that work with AI

The principle is simple: instead of hunting for AI use, you should design assignments where AI is either irrelevant, openly permitted, or impossible to hide. This is not theory — it is the way you already work with other aids.

Here are seven concrete assignment types that work. They are not mutually exclusive, and you do not have to implement them all. Pick one or two, and start there.

4.1 Process assignments with in-class checkpoints

Assignments where the components are both made *and* documented in class — not at home. The difference from ordinary staged submissions is that you have seen the parts come into being yourself, so AI cannot have written them.

Example — a Danish essay across four weeks:

- **Week 1:** 15-minute handwritten brainstorm in class, photographed and submitted the same day.
- **Week 2:** Two-minute oral presentation on the theme. You note the main points.
- **Week 3:** The outline is written in class — preferably on a locked school device or on paper.
- **Week 4:** Final essay submitted. The student answers 2–3 questions about their own text in a short oral follow-up (2–3 minutes).

Why it works: AI can help between checkpoints, but not during them. If the finished product does not match what you have seen along the way, it is obvious — without you having to trust a detector. Process-based assessment with observed component parts produces around 40 per cent fewer cases of cheating, according to *Frontiers in Education* (2024).

4.2 Locally grounded assignments

Assignments built on content the AI does not have access to: today's class discussion, a field trip, a specific copy of a book, your local area. When the answer requires specific details — quotations, page numbers, names, images — it becomes visible if the student has guessed or had the AI make something up.

Examples:

- "Describe three different positions from the discussion in class on Thursday 8 May. Quote at least one student (anonymously)."
- "Analyse three specific metaphors in Chapter 4 of [the book we read last week] with page and line references."
- "Take a photo of a building in your school district that demonstrates [an architectural/historical principle]. Explain the connection in 200 words."
- "Interview a local craftsperson (5 min). Extract three professional terms from the conversation and explain them."

Why it works: AI has broad knowledge but no knowledge of *your* class, today's debate, or the specific book on the student's shelf. UNESCO and the Cornell Center for Teaching Innovation point in 2025–26 to assignments rooted in a concrete situation as the most robust form of AI resistance.

4.3 Oral presentations with follow-up questions

The student presents orally and must be able to answer your questions.

Possible framing: *"You may use ChatGPT to research, but you have to be able to explain everything without notes. After the presentation, I will ask three random questions."*

Why it works: Even if students use AI to prepare, they have to *understand* the content to present it and answer questions. AI cannot deliver the presentation for them.

4.4 Multimodal assignments

Combine writing with video, drawing, models, experiments or other outputs the AI cannot produce.

Examples:

- A two-minute video explaining photosynthesis.
- An infographic about the Napoleonic Wars with a short written explanation.
- A physical model of a cell, with video documentation of the process.

Why it works: AI can write, but it cannot record a video of the student, build a physical model, or draw with the student's hand.

4.5 Critical evaluation of AI

Make AI part of the assignment itself — let the student assess the AI's quality rather than let the AI do the producing.

Example (history):

"Ask ChatGPT to write a text about [a historical event]. Analyse the text:

- *Which facts are correct/incorrect?*
- *Which perspectives are missing?*
- *Which sources should the AI have used?*
- *Rewrite the text so that it is more nuanced."*

Why it works: The student has to demonstrate critical thinking and factual knowledge in order to evaluate the AI. A side benefit is that the student notices where the AI gets things wrong.

4.6 Collaborative in-class assignments

Group assignments carried out during lesson time.

Structure:

- Groups of three or four students.

- All the work takes place in the classroom across the next three lessons.
- You can observe and ask questions along the way.
- Groups may use AI as a research tool, provided they document it.

Why it works: You see the work process. Cheating is impossible when the work happens in front of you.

4.7 Portfolio with rolling submissions

The student builds a portfolio across a whole unit with multiple work products.

Example (year 8, Danish):

The portfolio must include:

- Three different text types (op-ed, short story, analysis).
- A self-produced podcast episode (3–5 min).
- A reflection log after each lesson.
- Peer feedback the student has given to others.
- A final reflection on their own development.

Why it works: The variety and the rolling submissions make it very hard to cheat consistently. You see the student develop over time.

5. “I have 28 students — this will take forever”

This is the honest objection. Oral follow-ups, four staged submissions and portfolios all require time you may not have. Three practical adjustments ease the pressure:

- **Spot-check rather than total coverage.** You do not have to do an oral follow-up with everyone. Pick 4–5 students per unit at random. That is enough for the whole class to take the preparation seriously, and it is manageable for you.
- **Change the assignment type, not the workload.** A five-minute video assignment is not slower to grade than an essay — it is just different. Multimodal assignments do not necessarily add time.
- **Let students read each other’s work first.** Mutual peer feedback catches much of what you would otherwise have to address yourself. It also gives the students a subject-matter conversation about quality.

This is not a perfect system. It is an adjustment. Start with one subject or one unit. One assignment redesigned is worth more than seven half-finished.

6. The conversation when you suspect a student

If you have a specific suspicion, the conversation matters more than anything else. The purpose is to understand, not to punish. The following five-stage framework has been tested by Danish lower-secondary teachers and can be adapted to your context.

Stage 1 — a curious opening

"I would like to hear more about how you approached this assignment. Can you tell me about your process?"

The tone is not accusatory. It is genuine curiosity.

Stage 2 — ask the student to explain

"Can you elaborate on this passage here? How did you arrive at this conclusion?"

If the student struggles to explain, that is an indication — but they may simply be nervous. Stay open.

Stage 3 — direct but empathetic

"I notice that the language in this text is very different from what I usually see from you. I am not angry, but I am curious — did you get help from outside?"

Give the student a way out: *"It is fine to get help. But I need to know, so I can assess what you yourself can do."*

Stage 4 — a learning opportunity (if the student admits AI use)

"Thank you for being honest. Let us talk about how you can use AI to learn instead of replacing learning. What if we tried again together?"

The focus is forwards, not backwards. Learning before shame.

Stage 5 — if the student denies, and you are still in doubt

This is the hardest scenario. You have no proof — only suspicion. Three principles:

- **Assess what you can assess.** Ask the student for an oral follow-up: *"Explain this passage in your own words, or write a short paragraph here in class on the same topic."* What you observe there is something you can assess without accusation.
- **Go to leadership instead of standing alone.** Academic integrity is not the teacher's responsibility alone. If you have repeated suspicions, document them and pass the matter on — that is precisely why a school has leadership.
- **Let it go if you have nothing more than a hunch.** A single suspicious assignment is not worth damaging the relationship over. Adjust the next assignment instead, so that AI use becomes visible or irrelevant.

Remember: A false accusation is worse than a single missed case of cheating. You have the rest of the school year to adjust the assignment design.

7. The official Danish framework

Here is what you need to know as a Danish teacher or school leader about the official rules (as of May 2026).

For exams and assessments

- **Folkeskolen final examinations:** Generative AI (ChatGPT, Claude, Gemini and the like) is *prohibited* in all written and oral examinations. Submitting AI-generated content as one's own counts as cheating. Source: the Ministry's examination rules.
- **Upper-secondary examinations (STX, HHX, HTX, HF):** Same rules. Prohibited in all examinations.
- **Permitted aids:** Spell check, grammar check, read-aloud and translation (compensatory IT) — not generative AI. Pay attention to the grey zone: a bilingual student may use a translation tool, but not ChatGPT to rephrase entire sentences. Discuss this concretely with the student and their parents.
- **2026 pilot:** Selected upper-secondary schools are trialling whether students may use generative AI during the 1-hour preparation for the oral English exam. The presentation itself remains live, so the speaking ability is tested without AI.

In teaching

- **The Ministry's "Seven recommendations for the Folkeskole"** (June 2025) on the use of generative AI — worth reading for school leadership and subject teams.
- **"Guidance on the lawful use of AI"** (January 2026), covering the whole education sector.
- **The GDPR angle is decisive.** If student assignments contain personally identifiable information, they may not be entered into an AI tool without a data processing agreement. This is not only a legal matter — it is also about what the students learn about responsible data use. Check your school's list of approved suppliers before recommending a tool.
- **The school board** can adopt local rules in addition to the national ones — but cannot override the examination rules.

In short: Students may use AI to *learn* — explain concepts, brainstorm, practise. They may not use AI to *produce* the final product that is being assessed. That is the line you, as a teacher, have to help them navigate.

8. Template: local policy for AI use

A workable local policy addresses four things: when AI is permitted, when it is not, how the student documents their AI use, and what consequences breaches carry. The template below can be adapted by year group and subject.

1. When AI is permitted

- Brainstorming and idea generation.
- Explaining difficult concepts.
- Checking grammar and spelling.
- Researching topics, with source attribution.
- When the teacher explicitly permits it.

2. When AI is not permitted

- To write whole assignments for you.
- To solve maths problems without understanding the solution.
- For final exam-equivalent projects.
- When the assignment specifically says “without AI”.

3. Documenting AI use

If the student has used AI, they must write at the bottom of the assignment:

> *“I used ChatGPT for [specific task, e.g. ‘explaining photosynthesis’ or ‘finding synonyms’]. I understand the content and have written it in my own words.”*

4. Consequences of breaches

- **First time:** Conversation with the teacher. The assignment is redone (with support).
- **Repeated cases:** Contact with the parents. Possible time-limited suspension from using AI tools in school hours (but not from compensatory IT such as read-aloud or spell check, if the student needs it).
- **At examinations/final assessments:** Follows the school’s general rules on cheating.

9. Do and don't

A short checklist for daily use:

Do

- Design assignments where process counts.
- Allow AI use for brainstorming and research, provided it is documented.
- Have honest conversations about academic integrity.
- Teach students *how* to use AI responsibly — it is part of their professional formation.
- Assess orally where possible.
- Focus on understanding, not just the product.
- Be open about using AI yourself in your own preparation.

Don't

- Trust AI detectors blindly.
- Accuse students without a conversation.
- Ban AI outright — they will use it anyway.
- Punish curiosity ("I just wanted to try it").
- Set the same assignment type year after year.
- Ignore the problem and hope it goes away.
- Shame students publicly.

10. The most important thing to remember

AI cheating is a symptom of outdated assignment types and a fear-based approach to technology. The solution is not more control, but better assignment design and honest conversations.

The goal is not to catch students cheating. The goal is to teach them to think critically, solve problems, and use tools — including AI — responsibly.

In the future, AI will be a part of almost all work. Your job as a teacher is to prepare students for that reality — not to keep them away from it.

In 2026, 92 per cent of students use AI in one form or another. The question is no longer whether you can keep AI out of the classroom. The question is whether your assignments mean that the student still has to think.

This document is an action paper. For the full research overview behind the points on learning outcomes and cognitive offloading, see

Artificial Intelligence in Danish K–9 Schools — A research overview, May 2026

---

by Martin Hasbirk Brandt.