

## Copyright © 2026 Med Kharbach, PhD

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. You are free to share (copy and redistribute the material in any medium or format) and adapt (remix, transform, and build upon the material) under the following conditions: appropriate credit must be given, a link to the license must be provided, and any changes must be indicated. The material may not be used for commercial purposes. If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

[License details: creativecommons.org/licenses/by-nc-sa/4.0/](https://creativecommons.org/licenses/by-nc-sa/4.0/)

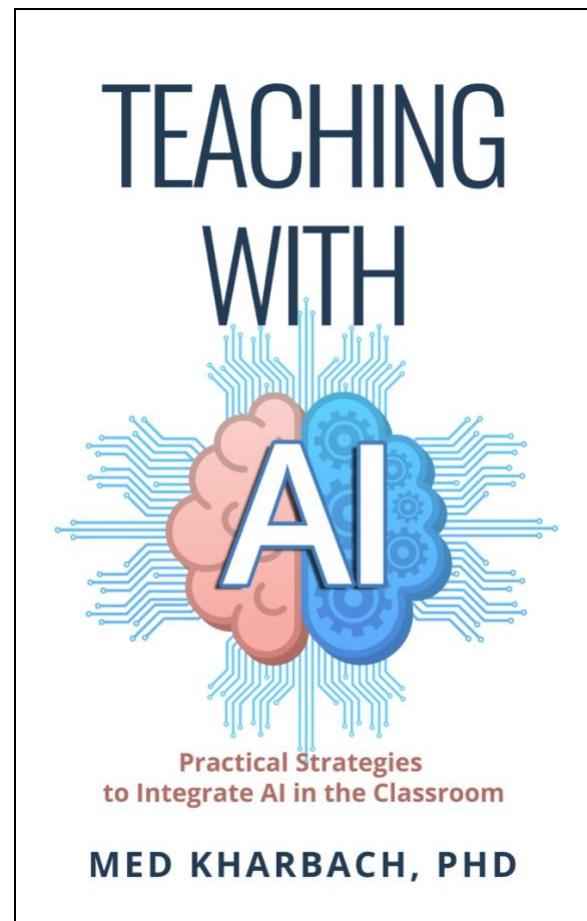
# Teaching with AI

## Practical Strategies to Integrate AI in The Classroom

In [Teaching with AI](#), I speak directly to you as a teacher working through real classroom questions about AI. The book helps you build strong AI literacy so you understand how these systems work, where they help, and where caution is needed.

I share concrete classroom strategies, examples drawn from practice, and ways to align AI use with sound pedagogy and professional judgment. My goal is to support you in using AI thoughtfully as part of your teaching, in ways that deepen learning and keep human expertise at the center.

Grab your copy [here](#).



## Introduction

Generative AI has created a problem that goes far deeper than cheating. When a tool like ChatGPT can write a coherent essay, solve a multi-step math problem, analyze a historical event, and produce a lab report, all in minutes, the entire notion of assessment comes into question. We built our assessment systems around the assumption that complex cognitive work was exclusively human. That assumption no longer holds, and we're now forced to ask a question many of us never expected to face: how do you measure genuine learning when a machine can produce the same or even better output?

Think about writing for a moment. For centuries, writing was the primary way we measured what students know. An essay, a research paper, a lab conclusion, a short answer response. Writing was the window into thinking. If a student could articulate a clear argument, organize evidence logically, and draw original conclusions, we took that as proof of learning. Now we have tools that can do all of this at a level that passes for competent student work, and often exceeds it. So when you assign a written assessment today, what exactly are you measuring? The student's ability to write? Their ability to think? Their ability to prompt an AI effectively? The line between those has blurred in ways that traditional assessment was never designed to handle. And it raises a question we should have been asking all along: were we ever really assessing thinking, or were we just assessing the ability to produce written text?

Here is what I believe is the deeper issue. The problem of assessment in the age of AI is, at its core, a problem of assessment literacy. I have mentioned this before in my earlier guide on [AI grading tools](#). Assessment literacy, the knowledge of how to design assessments that genuinely measure learning, is something rarely covered in teacher education programs (Popham, 2018). Most of us, when we started teaching a decade or two ago, walked into our classrooms with no specialized training in assessment

design. We had subject knowledge. We had classroom management strategies. We had curriculum frameworks. But assessment? That was trial and error. We copied what our own teachers did, followed department templates, and figured it out as we went. That approach worked well enough when students couldn't outsource cognitive work to a machine. It doesn't work anymore.

And here is where the conversation often goes wrong. Too many teachers and administrators blame students for using AI. The accusations fly: students are lazy, they're cheating, they don't want to learn. But the blame should be placed squarely on assessment design, not only on the students. We know they'll use AI whether we like it or not. [88% of UK university students](#) used generative AI for assessments in 2025, up from 53% just one year earlier (Freeman, 2025). The numbers in K-12 are climbing fast. You can ban AI, threaten consequences, install detection software, and students will still find ways around it. The question we need to ask is how to design assessments where using AI doesn't allow a student to bypass the actual learning.

In this guide, I share insights I've learned from fellow teachers, education researchers, and assessment specialists about redesigning assessments for the AI era. I'll walk through what the current research tells us, introduce some practical frameworks for rethinking your approach, and then get into specific strategies that teachers and researchers report as genuinely effective. I'll also talk about what doesn't work, and I'll close with creative ideas from educators who are finding their way through this same challenge.

## What the Research Says

The research on AI and assessment has been growing fast, and the findings are sobering.

ChatGPT can already answer 65.8% of exam questions correctly across 50 diverse university courses. That's the headline number from a [2024 PNAS study](#) (Borges et al., 2024), and in engineering the success rate climbed even higher for standard problem sets. Traditional exams are vulnerable precisely because they tend to test recall and standard application, the kinds of tasks AI handles best. But there's an encouraging flip side. AI performance dropped significantly on higher-order Bloom's taxonomy tasks: analyze, evaluate, create. That tells us where the vulnerability is, and where the opportunity lies.

Hardie et al [\(2024\)](#) tested 17 different assessment types against generative AI. Most of them crumbled. Standard essays, reports, and problem sets were the weakest links. The formats that actually held up? Audience-tailored assessments, observation by learner, and reflection on work practice. They share a common feature: each one requires something specific, personal, or situated that AI can't easily fake.

Here's a nuance that caught my attention. Research published in the [British Journal of Educational Technology \(2025\)](#) by Kofinas et al. found that even authentic assessments, the kind many institutions are rushing to adopt, don't automatically safeguard academic integrity. Markers in the study generated both false positives (accusing students who didn't use AI) and false negatives (missing students who did). No single assessment approach is enough on its own. You need layers.

And on the detection front? The evidence is even bleaker. Gaines [\(2025\)](#) reported that AI detection tools are unreliable, yet teachers keep using them anyway. OpenAI shut down its own AI text classifier because accuracy was too low. Vanderbilt, UT Austin, Montclair State, and Northwestern all advised faculty not to rely on Turnitin's

AI detection. The University of Maryland found no publicly available detector sufficiently reliable for institutional use. Marc Watkins, [writing on Substack](#), documented how unreliable detection actively harms students, with non-native English speakers disproportionately flagged.

As [Emma Whitford](#) (2025) put it: you can't AI-proof the classroom. But you can get creative.

## Frameworks Worth Mentioning

Before jumping into specific strategies, it helps to have a broader way of thinking about AI and assessment. I've selected three frameworks that educators and researchers have found particularly useful, and each one gives you a structured approach to deciding when and how to include or exclude AI from your assessments.

### 1. The AI Assessment Scale (AIAS)

Developed by Mike Perkins, Leon Furze, Jasper Roe, and Jason MacVaugh, the [AI Assessment Scale](#) has been adopted by hundreds of schools worldwide and translated into 30+ languages. Five levels of AI use, from zero to full creative exploration. Level 1 means no AI at all, completed in controlled environments. Level 2 allows AI for brainstorming and outlining only. Level 3 permits collaboration where students critically evaluate and revise AI output. Levels 4 and 5 open the door to full AI use and creative AI exploration. The real value is transparency. You specify on each

assignment exactly what level is permitted, students know where they stand, and the ambiguity that fuels so many "is this cheating?" conversations just evaporates.

## The AI Assessment Scale

<b>1</b>	<b>NO AI</b>	The assessment is completed entirely without AI assistance in a controlled environment, ensuring that students rely solely on their existing knowledge, understanding, and skills <b>You must not use AI at any point during the assessment. You must demonstrate your core skills and knowledge.</b>
<b>2</b>	<b>AI PLANNING</b>	AI may be used for pre-task activities such as brainstorming, outlining and initial research. This level focuses on the effective use of AI for planning, synthesis, and ideation, but assessments should emphasise the ability to develop and refine these ideas independently.  <b>You may use AI for planning, idea development, and research. Your final submission should show how you have developed and refined these ideas.</b>
<b>3</b>	<b>AI COLLABORATION</b>	AI may be used to help complete the task, including idea generation, drafting, feedback, and refinement. Students should critically evaluate and modify the AI suggested outputs, demonstrating their understanding.  <b>You may use AI to assist with specific tasks such as drafting text, refining and evaluating your work. You must critically evaluate and modify any AI-generated content you use.</b>
<b>4</b>	<b>FULL AI</b>	AI may be used to complete any elements of the task, with students directing AI to achieve the assessment goals. Assessments at this level may also require engagement with AI to achieve goals and solve problems.  <b>You may use AI extensively throughout your work either as you wish, or as specifically directed in your assessment. Focus on directing AI to achieve your goals while demonstrating your critical thinking.</b>
<b>5</b>	<b>AI EXPLORATION</b>	AI is used creatively to enhance problem-solving, generate novel insights, or develop innovative solutions to solve problems. Students and educators co-design assessments to explore unique AI applications within the field of study.  <b>You should use AI creatively to solve the task, potentially co-designing new approaches with your instructor.</b>



Perkins, Furze, Roe & MacVaugh (2024). The AI Assessment Scale

## 2. Danny Liu's Two-Lane Approach

Danny Liu at the University of Sydney has a different take. He proposes [splitting assessments into two lanes](#). Lane 1 focuses on verifying what students actually know: proctored tests, oral exams, in-class work with no AI access. Lane 2 goes the other direction entirely, asking students to demonstrate how they use AI critically and thoughtfully within their discipline. Liu's core argument resonates with me. Trying to outrun or 'outdesign' AI with increasingly complex assignments is a losing game. A

better move is accepting that both lanes serve learning, and designing your course with some of each.

## The FACT Framework

Published in [Frontiers in Education \(2025\)](#) by Elshall and Badir (2025), the FACT framework (Framework for AI-Conscious Teaching) started in environmental data science but applies broadly. The core idea: identify which learning outcomes require human-only demonstration, which benefit from AI collaboration, and build your assessments around those answers. It gives you a structured decision-making process, which is especially helpful when you're staring at a syllabus full of assignments and wondering which ones need rethinking.

## Assessment Strategies

Here are the strategies that research and teacher experience point to as genuinely effective. My recommendation: don't try to use all of them at once. Pick two or three that fit your context and combine them.

### 1

## Oral Assessments and Live Defense

This is the one researchers and teachers agree on most strongly. For instance, Hartman (2025) [research](#) confirmed what many of us already suspected: oral exams are naturally AI-resistant. Students have to think on their feet, respond to follow-up questions, and demonstrate understanding in real time. No amount of prompting ChatGPT can prepare you for a teacher who asks "OK, but why did you choose that particular approach?"

The formats vary. You might use 15-minute oral exam slots with one straightforward and one probing question. Or viva voce sessions where students defend their written submissions. Or Q&A rounds after presentations. Some teachers have students record

podcast-style audio where they explain a concept as if teaching it. Others ask for video walkthroughs of their reasoning.

Time is the honest concern. Fifty students at 15 minutes each means 12+ hours. The workaround many teachers use: grade the essay first, then have a 5-minute conversation about it. That shorter window still reveals whether the student actually wrote what they submitted.

## 2

## Process-Based Assessment

If you can only adopt one strategy from this whole guide, make it this one.

Grade the process, not just the product. Break larger assignments into stages: research proposal, outline, rough draft, revision, final submission. Grade each one separately. Have students work in Google Docs or Word Online where the revision history is visible. Ask for research journals documenting how they found and evaluated their sources. Request track-changes documents showing their revisions between drafts, with written explanations for each choice.

Why does this work so well? AI-generated text shows up fully formed. One editing session, one paste, done. There's no messy brainstorming, no wrong turns, no gradual development of ideas. When you pull up a student's revision history and see an entire 1,500-word essay materialize in a single moment, the story tells itself. Human writing leaves footprints. AI writing appears out of thin air.

### 3

## Personal Experience and Authentic Reflection

AI produces competent analytical writing all day long. What it can't do is remember being confused during your Tuesday lab. [Dr. Catlin Tucker recommends](#) pushing students to explain their own thinking process, their decision-making, their reasoning, in ways that demand real metacognitive awareness.

Some prompts that work: "Which step in [topic] confused you most, and what analogy did YOU create to understand it?" Or: "Describe a moment during our class this week when your prediction turned out wrong. What did you learn from that?" Or: "How has your understanding of this topic changed since the beginning of the unit? Point to specific class discussions or readings that shifted your thinking."

AI has no classroom memories. It didn't attend your lab. It wasn't there for the discussion that went sideways on Thursday. That's exactly the gap you can design around.

### 4

## Local and Community Context

Ground your assignments in knowledge that AI simply doesn't have. The classic history prompt "What were the causes of World War I?" is exactly the kind of question ChatGPT handles comfortably. But try this: "Based on our class visit to the local war memorial, explain how communities in our region experienced the impact of global conflicts." That requires specific, firsthand knowledge no AI model possesses.

Other examples: research projects built on local environmental data, case studies about decisions your city council actually made, assignments that reference a specific

guest speaker who visited your class, or essays drawing on a shared field trip experience. The more local and specific you go, the less useful AI becomes. AI knows everything in general. It knows nothing about your school.

## 5

## Collaborative and Group Assessments

AI works alone. That's its fundamental limitation in this context. Students working in genuine collaboration, negotiating roles, debating approaches, building on each other's ideas in real time, produce something AI can't replicate.

The key word is genuine. A group project where one person does all the work is just as vulnerable as an individual assignment. The collaboration needs to be visible and accountable. Have students keep decision logs about how they divided responsibilities and worked through disagreements. Build in individual reflection components. Conduct brief group interviews where each member explains their contribution. The documentation is what makes group work AI-resistant, not the group work itself.

## 6

## Alternative and Multimodal Formats

Text is AI's strongest medium. Move away from it.

Video explanations where students teach a concept. Podcast episodes weaving research with personal narrative. Flowcharts and concept maps. Physical models.

Annotated artwork. Comic strips that walk through a scientific process. These formats demand a kind of authentic expression that a text prompt and a chatbot can't produce.

One favorite from teachers I've read about: require a 3-minute video explanation of the essay before you'll grade it. If the student wrote it, they'll talk about it naturally. If AI wrote it, you'll hear the disconnect within the first 30 seconds.

## 7

## In-Class and Timed Assessments

The blue book has made a comeback. Gaines [\(2026\)](#) reported on a high school English teacher who went fully analog after watching AI change her students' relationship to writing. In-class essays, timed exams with unknown questions, quick low-stakes quizzes, and live problem-solving sessions all keep AI out of the equation because the devices stay out of the room.

Straightforward? Yes. But the trade-offs are real. Timed writing favors fast processors and disadvantages students with test anxiety, learning disabilities, or language barriers. If you go this route, pair it with other assessment types so every student has multiple ways to show what they know. A timed essay alone is not a complete picture of anyone.

## 8

## Higher-Order Thinking Prompts

Borges et al. (2024) found that AI performs significantly better on lower Bloom's taxonomy levels (remember, understand, apply) and struggles at the top (analyze, evaluate, create). Your prompts should aim high.

Ask students to critique an AI-generated response and identify where it falls apart. Build multi-step case studies where the "right" answer depends on context and professional judgment. Create ethical dilemmas with no clean solution. Ask students to synthesize conflicting sources and make their own argument.

Here's one I particularly like. Give students an AI-generated response to your assignment prompt and make the evaluation the assignment. What did the AI get right? Where is it shallow? What would a student who actually understands this material add that the AI missed? You're building critical AI literacy and testing subject knowledge with the same prompt.

## Practical Tips

Some tips drawn from teacher experience, online discussions, and the research.

**Start with your learning objectives, not the AI threat.** What do you actually need students to demonstrate? If the answer is "clear analytical writing," you need a strategy that verifies the student produced the writing. If the answer is "understanding of [topic]," maybe a written essay isn't the only path. Design backward from what matters most.

**Use the ChatGPT test.** Before you assign anything, paste your prompt into ChatGPT. If the AI produces a passable response, your prompt needs reworking. Keep redesigning until the AI output is clearly incomplete or generic compared to what a student who actually did the thinking would produce. Five minutes of testing saves you weeks of frustration.

**Layer your strategies.** No single approach is enough. The teachers who report the best results combine two or three: process documentation plus a personal context component, or a written assignment followed by an oral defense, or group collaboration plus individual reflection. Each layer closes a gap that another one leaves open.

**Build transparency into your approach.** Tell students exactly what AI use is permitted and what isn't, ideally using a framework like the AI Assessment Scale. Students are far less likely to cross a line they can actually see. And when you're open about your reasoning, the conversation shifts from surveillance to shared understanding about what genuine learning looks like in your course.

**Have the conversation with students.** Some teachers allow AI on certain assignments, with one condition: students must document exactly how they used it. What prompts did they give? What did they keep, revise, or discard? Why? This turns AI use into a learning opportunity. It also builds exactly the kind of critical thinking about AI tools that students need.

**Consider equity at every step.** Oral exams create anxiety for some students. Video submissions require tech access and camera comfort. Timed writing disadvantages students with processing differences. Every AI-resistant strategy has equity implications, and the honest response is to build in choice. Give students multiple pathways to demonstrate their learning, and be transparent about why you're offering them.

**Start small.** You don't need to overhaul everything this semester. Pick one assignment and add one AI-resistant element: a process component, a brief oral follow-up, or a personal reflection requirement. See what happens. Adjust. Then expand from there.

**The "Code of Conduct" approach.** Some teachers co-create an AI use agreement with their class at the start of the term. Together they discuss what AI use is appropriate, what crosses a line, and why those boundaries matter. The social accountability this generates is surprisingly powerful. Students who help define the norms tend to be far more invested in upholding them.

**Video explanation requirements.** A growing number of teachers ask students to record a short video (3-5 minutes) explaining their submitted work before it receives a grade. It doesn't need to be polished. It just needs to show the student can talk fluently about what they wrote, why they made specific choices, and what they'd change given more time. Simple addition. Huge payoff.

**Peer and AI Review + Reflection (PAIRR).** A [model published in ScienceDirect](#) by Sperber et al. (2025) where students first get AI-generated feedback on their drafts, then get peer feedback, and then write a reflection comparing the two. Which feedback was more helpful? Where did the AI miss something the peer caught? Where was the AI actually more useful? It builds metacognitive skills and teaches students to evaluate feedback critically, regardless of its source.

**Portfolio assessment across a term.** Students document their entire learning journey: brainstorms, rough drafts, feedback received, revision decisions, final products, and reflective narratives about their growth. The portfolio format makes sustained AI use much harder to pull off because the evidence of learning is distributed across weeks of documented work, not packed into one final submission.

## Conclusion

The assessment challenge AI has created is real. But in many ways, it's also an overdue invitation to rethink how we measure learning. Many of the practices AI has disrupted were already limited. The five-paragraph essay tested formula-following as much as critical thinking. The recall-heavy exam measured memorization as much as understanding. AI just made the cracks impossible to ignore.

Every strategy in this guide shares a common thread: they ask students to show their thinking, not just their output. A staged writing process, an oral conversation about their work, a video walkthrough, a local research project, a reflection grounded in personal experience. These approaches push assessment toward what it was always

supposed to measure: genuine understanding, real growth, the ability to think through complex problems with your own mind.

You don't need to adopt all eight strategies tomorrow. Start with the one that fits your current workflow. Add a process component to an existing assignment. Try a brief oral follow-up after a written submission. Build in a personal reflection prompt that anchors the assessment in something only that particular student could write. Small changes, applied consistently, add up to assessments that are far harder for AI to shortcut and far better at showing what your students actually know.

## About the Author

Med Kharbach, PhD, is an educator, researcher, and the editor of Educators Technology ([educatorstechnology.com](https://www.educatorstechnology.com)). A former K-12 teacher with over a decade of classroom experience, Med currently serves as part-time faculty at Mount Saint Vincent University. He is the author of *Teaching with AI: Practical Strategies to Integrate AI in The Classroom*, and co-author of *The AI Turn in Academic Research* with Dr. Jonathan Woodworth. His work focuses on the intersection of technology, pedagogy, and practical classroom applications. You can read more about Med and his research on his professional website [www.medkharbach.com](https://www.medkharbach.com).

## References

- Borges, B., Foroutan, N., Bayazit, D., & EPFL Data Consortium. (2024). Could ChatGPT get an engineering degree? Evaluating higher education vulnerability to AI assistants. *Proceedings of the National Academy of Sciences*, 121(49), e2414955121. <https://doi.org/10.1073/pnas.2414955121>
- Duke University (2025). "Authentic Assessment over Surveillance." <https://lile.duke.edu/blog/2025/10/authentic-assessment-over-surveillance/>
- Elshall, A. S., & Badir, A. (2025). Balancing AI-assisted learning and traditional assessment: The FACT assessment in environmental data science education. *Frontiers in Education*, 10, 1596462. <https://doi.org/10.3389/feduc.2025.1596462>
- Freeman, J. (2025, February). Student generative AI survey 2025 (HEPI Policy Note 61). *Higher Education Policy Institute*. <https://www.hepi.ac.uk/wp-content/uploads/2025/02/HEPI-Kortext-Student-Generative-AI-Survey-2025.pdf>
- Gaines, L. V. (2025, December 16). Teachers are using software to see if students used AI. What happens when it's wrong? *NPR*. <https://www.npr.org/2025/12/16/nx-s1-5492397/ai-schools-teachers-students>
- Gaines, L. V. (2026, January 28). To keep AI out of her classroom, this high school English teacher went analog. *NPR*. <https://www.npr.org/2026/01/28/nx-s1-5631779/ai-schools-teachers-students>
- Hardie, L., Lowe, J., Pride, M., Waugh, K., Hauck, M., Ryan, F., ... & Richardson, H. (2024). Developing robust assessment in the light of Generative AI developments. *NCFE; The Open University*. <https://oro.open.ac.uk/99447/>
- Hartmann, C. (2025). Oral exams for a generative AI world: Managing concerns and logistics for undergraduate humanities instruction. *College Teaching*. Advance online publication. <https://doi.org/10.1080/87567555.2025.2558563>
- Kofinas, A. K., Tsay, C. H.-H., & Pike, D. (2025). The impact of generative AI on academic integrity of authentic assessments within a higher education context. *British Journal of Educational Technology*. Advance online publication. <https://doi.org/10.1111/bjet.13585>

Liu, D. (2023, July 4). What to do about assessments if we can't out-design or out-run AI? *LinkedIn*. <https://www.linkedin.com/pulse/responding-generative-ai-assessments-semester-2-2023-danny-liu/>

Perkins, M., Roe, J., & Furze, L. (2024). The AI Assessment Scale revisited: A framework for educational assessment (Preprint). December 2024. <https://arxiv.org/abs/2412.09029>

Popham, W. J. (2018). *Assessment literacy for educators in a hurry*. ASCD.

Sperber, L., MacArthur, M., Minnillo, S., Stillman, N., & Whithaus, C. (2025). Peer and AI Review + Reflection (PAIRR): A human-centered approach to formative assessment. *Computers and Composition*, 76, 102921. <https://doi.org/10.1016/j.compcom.2025.102921>

Tucker, C. (2024). "5 Tips for Designing AI-Resistant Tasks." <https://catlintucker.com/2024/10/ai-resistant-tasks/>

Walton Family Foundation. (2025, June 25). The AI dividend: New survey shows AI is helping teachers reclaim valuable time. <https://www.waltonfamilyfoundation.org/>

Watkins, M. (2023, September 3). Beyond ineffective: How unreliable AI detection actively harms students. *Substack*. <https://marcwatkins.substack.com/p/beyond-ineffective-how-unreliable>

Whitford, E. (2025, December 16). You can't AI-proof the classroom, experts say. Get creative instead. *Inside Higher Ed*. <https://www.insidehighered.com/news/faculty-issues/learning-assessment/2025/12/16/you-cant-ai-proof-classroom-experts-say-get>