
Critical Thinking Activities for the Age of AI

Hands-On Exercises for Faculty and Students

Dr. Med Kharbach, PhD

educatorstechnology.com

2026



Creative Commons License

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

You are free to share and adapt this material for non-commercial purposes, as long as you give appropriate credit and distribute under the same license.

© 2026 Med Kharbach

Introduction

A growing body of research is starting to map the cognitive impact of generative AI on how our students think. I've reviewed much of this academic work on my blog at medkharbach.com/blog, and I won't bog you down with all of it here. But if you want to go down that rabbit hole, some of the most significant studies include Gerlich (2025) on cognitive offloading and its effect on critical thinking, Fan et al. (2025) on metacognitive laziness, Shaw and Nave (2026) on what they call cognitive surrender, Kosmyna et al. (2025) who used brain imaging to measure the neurological effects of using ChatGPT for writing tasks, Barcaui (2025) on how ChatGPT use reduces knowledge retention, Bastani et al. (2025) who found that generative AI without guardrails actively harms learning in mathematics, and Sourati et al. (2026) on how large language models are homogenizing the way we think and express ourselves.

The themes running through this research are concerning. Students who rely on AI for cognitive tasks show reduced critical thinking engagement, diminished metacognitive monitoring, and a tendency to accept AI-generated content without questioning it. Some researchers describe this as a kind of thinking atrophy: the mental muscles that students need for deep reasoning start to weaken when AI does the heavy lifting for them. The brain imaging data from Kosmyna et al. is particularly striking, showing measurably lower cognitive effort during AI-assisted tasks compared to independent work. And it's not just about thinking. Sourati et al. and Abdulhai et al. (2026) show that AI is also flattening the diversity of how we write and communicate, pushing everyone toward the same patterns, the same structures, the same voice.

I take these findings seriously. And as educators, we should all be paying attention. But I also have two reservations about this body of research that I think are worth naming.

First, representativeness. The populations in many of these studies are narrow. Factors like institutional context, discipline, prior AI literacy, age, and how much experience students have had with AI tools all differ across studies. A finding from a sample of psychology undergraduates at a single university doesn't automatically tell us what's happening in a nursing program across the country or in a hospitality management class in a completely different educational culture. The patterns are suggestive but definitely not universal.

Second, time. Generative AI is still a very new technology and measuring its cognitive impact in a valid and rigorous way takes years. The strongest evidence will come from longitudinal studies that track participants over extended periods, and we simply don't have those yet. Most of what we're seeing are snapshots: cross-sectional studies, short-term experiments, self-report surveys. They point in a direction, and that direction is worth taking seriously, but the definitive answers will need a few more years to arrive.

That said, the cognitive concerns these studies raise are valid. I'm not dismissing them. I'm asking for patience and nuance in how we interpret them, and I'm saying that we shouldn't wait for perfect evidence before we act.

Now, the question that probably brought you to this guide in the first place: what do we actually do about it?

I've been thinking about this a lot, and here's where I land. The way our students use AI is what counts, not the simple fact that they use it. Right now, many students approach AI as an answer machine. You type a prompt, wait for the output, copy it, and move on. That's the pattern that produces cognitive offloading and metacognitive laziness. That's where the thinking stops.

But what if we designed activities that pushed students to engage with AI differently? What if the AI interaction itself became the site of critical thinking, not the shortcut around it? What if students had to challenge AI output, question its assumptions, verify its claims, spot its reasoning gaps, and compare its responses against their own knowledge and judgment?

That's the purpose of this guide. It grew out of a workshop I delivered at the NENA ICHRIE 2026 conference titled "AI in the Classroom: Enabling Critical Thinking." For that workshop, I compiled a collection of hands-on activities drawn from educators, researchers, and practitioners working at the intersection of AI and pedagogy. I adapted each activity to connect directly to Robert Ennis's (2015) framework of critical thinking skills, which identifies six core competencies: analyzing arguments, judging source credibility, handling ambiguity, evaluating assumptions, recognizing fallacies, and asking clarification questions.

The activities in this guide are practical, flexible, and designed to work across disciplines and student levels. Some take ten minutes. Others need a full class session. You can use them as standalone exercises, string several together into a workshop sequence, or embed them into your existing course design. Each activity targets specific skills from Ennis's framework, so you can match activities to the thinking competencies you most want to develop in your students.

The goal is simple: don't ban AI, don't fear it, and don't let students sleepwalk through it. Use it as a training ground for the kind of rigorous, skeptical, engaged thinking that every discipline needs.

Critical Thinking Activities

The thirteen activities in this section are designed to turn AI interactions into critical thinking exercises. And as I mentioned earlier, each one targets specific skills from Robert Ennis's (2015) framework: analyzing arguments, judging source credibility, handling ambiguity, evaluating assumptions, recognizing fallacies, and asking clarification questions.

Some activities focus on two or three of these skills. A few touch all six.

I've drawn these activities from a range of sources: fellow educators, published research, open teaching repositories, a TED talk, and my own classroom experience.

For each activity, I've included a short narrative explaining where it comes from and why I think it works, followed by a summary table with timing, the CT skills it targets, step-by-step instructions, and the original source. You'll notice that most activities take between ten and twenty minutes. A few need closer to thirty. All of them can be adapted to different disciplines, adjusted for different student levels, and run with whatever AI tool your institution supports.

One more thing before you start. The activities themselves are solid, and I'm confident they'll generate productive engagement in your classroom. But the activity is only half the work. The real payoff comes from what happens afterward: the debrief, the reflective memo, the metacognitive checkpoint, whatever form you give it.

Each time you run one of these activities, have your students write a short reflective memo at the end. Ask them to think through what they noticed about the AI's output, where their own assumptions were challenged, what surprised them, what they'd do differently if they ran the same prompt again, and how the exercise changed the way

they think about using AI in their coursework. That reflection is where the critical thinking actually takes root. The activity opens the door. The reflective writing is where students walk back through it and make sense of what they experienced.

Don't skip this step. A well-run activity with no reflection produces a good classroom moment. A well-run activity followed by a reflective memo produces a thinking habit that students carry with them long after the class ends.

Activity 1: Push-Back Protocol

I first came across this activity through Tawnya Means, who designed it as a structured way to get students to challenge AI output at face value. I adapted it here to work with three of Ennis's critical thinking skills: analyzing arguments, evaluating assumptions, and asking clarification questions.

The setup is straightforward. Students prompt an AI tool with a question relevant to their discipline. Could be anything: a policy question, a case analysis, a historical interpretation. Then they read the response and identify the main claims being made. That's where the real work begins. They go through multiple rounds of structured questioning. They demand evidence for claims. They ask what's being assumed. They look for alternative perspectives the AI didn't mention. They stress-test the reasoning behind each point.

Each round maps to a specific skill on the rubric, so what starts as a routine chatbot exchange becomes an active critical thinking exercise. What I like about this activity is its simplicity. It works with any discipline, any student level, and any AI tool. And it teaches a habit that transfers: don't accept the first answer. Ask the next question.

Time	15-20 min
CT Skills	Analyze arguments, Evaluate assumptions, Ask clarification questions
How to Run It	<ol style="list-style-type: none"> 1. Prompt AI with a question relevant to your discipline. 2. Read the response and identify the main claims being made. 3. Challenge the response: ask for evidence, sources, and reasoning behind each claim. 4. Question assumptions: what is being taken for granted in the response? 5. Seek alternatives: ask AI if there are opposing views or exceptions.
Source	<p><i>Means, T. (2025, December 22). The push-back protocol: Teaching students to challenge AI, not accept it. The Collaboration Chronicle: Human+AI in Education.</i></p> <p>https://tawnyameans.substack.com/p/the-push-back-protocol-teaching-students</p>

Activity 2: Hallucination Detective

I found this activity through [AI for Education](#), originally designed by Michael Kingston. I adapted it here to zero in on Ennis's skill of judging source credibility.

Students receive an AI-generated text with errors built into it: fake citations, wrong dates, made-up statistics, and confident-sounding claims with no factual basis. Their job is to play detective. They go through the text line by line, flag what looks suspicious, and verify claims against reliable databases and sources.

You can open with the [Steven Schwartz case](#), the New York lawyer who submitted a ChatGPT-generated legal brief stuffed with entirely fabricated case citations and was sanctioned by a federal judge. That story sets the stakes fast: AI hallucinations aren't a minor inconvenience. They've already caused real professional consequences.

But the real learning comes in the debrief. Why does AI hallucinate? What makes hallucinated text so convincing? And what verification habits can students build into

their workflow to catch it before it causes damage? Those are the questions that stick.

Time	15-20 min
CT Skills	Judge source credibility, Recognize fallacies
How to Run It	<ol style="list-style-type: none"> 1. Present the real-world case of a lawyer whose ChatGPT-generated legal brief contained entirely fabricated case citations. 2. Distribute AI-generated text samples with known hallucinations embedded. 3. Students fact-check each claim, citation, and statistic using reliable databases. 4. Groups present findings: what was fabricated and how they verified it. 5. Discuss: why does AI hallucinate and what strategies can guard against it?
Source	<p>Kingston, M. (2023). <i>Hallucination detective [Lesson plan]. AI for Education.</i> https://www.aiforeducation.io/curriculum/lesson-3</p>

Activity 3: Human vs. AI Blind Comparison

This is one I created for a graduate course I taught at MSVU, and it's become one of my go-to activities because the setup is so minimal and the conversation it generates is always rich. You take two essays on the same topic, one written by a human and one generated by AI, and present them side by side with no labels. Students read both and try to figure out which is which.

The guessing part is fun, but it's not the point. The point is what students notice when they're forced to read closely and compare. They start picking up on patterns: the AI text sounds polished but stays vague, sources are missing or hard to trace, and the reasoning rarely goes below the surface. The human text might be rougher in places, but it tends to take positions, cite specific evidence, and show the kind of unevenness that comes from actually thinking through a problem.

That conversation connects directly to several of Ennis's skills, especially judging source credibility, analyzing arguments, and handling ambiguity. And it teaches

students something important about AI-generated writing: fluency isn't the same as accuracy. A paragraph can read beautifully and say almost nothing. You can run this in any discipline with minimal preparation, and it scales easily from a ten-minute warm-up to a full-class discussion.

Time	15-20 min
CT Skills	Analyze arguments, Judge source credibility, Handle ambiguity
How to Run It	<ol style="list-style-type: none"> 1. Prepare two essays on the same topic (one human-written, one AI-generated) without labels. 2. Participants read both texts carefully and take notes. 3. Individually, each person identifies which they believe is AI-generated and documents their reasoning. 4. Small group discussion: compare observations and rationale. 5. Reveal the answer. Discuss: what patterns distinguish AI writing from human writing?
Source	<i>This is an activity I created for one of the courses I taught.</i>

Activity 4: AI Output Audit with Rubric

This activity uses the Critical Thinking Rubric included in the guide, which I built based on Ennis's (2015) six critical thinking skills and inspired by Katy McMahon's course on Critical Thinking in the Age of AI through Maven Analytics. It's one of the few activities that touches all six skills at once, and the structure is straightforward enough to work in any discipline.

Students receive a piece of AI-generated text and score it across all six standards: argument analysis, source credibility, ambiguity, assumptions, fallacies, and clarification. They use a simple rubric scale: Weak (1), Moderate (2), Strong (3) for each standard. Everyone scores individually first, then moves into pairs or small groups to compare.

The most productive moments come from the disagreements. When two people score the same text differently, they have to articulate why, and that conversation is critical thinking in action. One student might rate the argument analysis as strong because the text sounds logical. Another might rate it weak because no evidence backs the claims. Both have to defend their reasoning, and both walk away with a sharper eye for what counts as a strong argument versus what just sounds like one.

You can use this with any AI-generated content in any discipline. I've found it works best when the text is on a topic students already know something about, so they can bring their own knowledge into the evaluation.

Time	15-20 min
CT Skills	All six critical thinking skills
How to Run It	<ol style="list-style-type: none">1. Distribute AI-generated text on a topic relevant to your discipline.2. Each participant scores the text using the rubric (Weak 1, Moderate 2, Strong 3) for each standard.3. In pairs or small groups, compare scores and discuss where you disagree.4. Identify which standards the AI performed best and worst on.5. Discuss: what does this reveal about AI strengths and blind spots?
Source	<i>McMahon, K. (2025). Critical thinking in the age of AI [Online course]. Maven Analytics.</i> https://mavenanalytics.io/course/critical-thinking-in-the-age-of-ai

Activity 5: Text Annotation Exercise

This comes from the WAC Clearinghouse, where it's framed as a way to use AI-generated text as a prompt for critical analysis. The original assignment was developed by J. P. Pardo-Guerra as part of the TextGenEd collection. I adapted it here to connect with Ennis's skills of analyzing arguments, evaluating assumptions, and recognizing fallacies.

Students receive an AI-generated essay or report and go through it with a pen, or digitally with a commenting tool like Hypothesis, marking everything that raises a flag: factual errors, unsupported claims, missing context, hidden assumptions. Then they organize their findings into a simple three-column table: what the AI said, what's wrong with it, and how to fix it.

The table format is what makes this activity work. You can't just say "this feels off." You have to name the problem and propose a fix. That forces a level of precision that most classroom discussions don't reach on their own. Students have to move beyond gut reactions and into specific, evidence-based critique. And when groups share their tables at the end, the patterns become visible: where does AI tend to go wrong? What kinds of errors repeat across different topics? Those are the questions that build lasting analytical habits.

Time	20-25 min
CT Skills	Analyze arguments, Evaluate assumptions, Recognize fallacies
How to Run It	<ol style="list-style-type: none"> 1. Provide an AI-generated essay or report on a course-relevant topic. 2. Participants annotate the text by highlighting issues in categories: factual errors, unsupported claims, missing context, hidden assumptions. 3. For each issue, write a brief note explaining what is wrong and how to fix it. 4. Create a summary table with three columns: Original Text, Problem, Proposed Fix. 5. Share findings with the group and identify the most common issues.
Source	<p><i>Pardo-Guerra, J. P. (n.d). Using AI text as prompts for critical analysis. In TextGenEd collection. WAC Clearinghouse.</i></p> <p>https://wacclearinghouse.org/repository/collections/textgened/rhetorical-engagements/using-ai-text-as-prompts-for-critical-analysis/</p>

Activity 6: Google vs. ChatGPT Side-by-Side

I adapted this one from a piece by Lila Shroff in *The Atlantic* about people outsourcing their thinking to AI. The article profiles individuals who've started relying on ChatGPT for decisions they used to make on their own, and it got me thinking about how students interact with search tools versus chatbots.

This one takes two minutes to set up and always sparks a good conversation. Students pick a factual question in their field and search it on both Google and ChatGPT. Google returns a list of links, each with a source you can evaluate and cross-reference. ChatGPT usually returns one confident, polished paragraph with no visible sources. The contrast is immediate.

Students start asking: which answer would you trust and why? Which one actually gives you the tools to verify what you're reading? This connects directly to Ennis's skills of judging source credibility and handling ambiguity. I also use it to bring back the fluency trap idea: just because something reads well doesn't mean it's accurate. A well-constructed sentence with no citation is still an unsupported claim, no matter how confident it sounds.

Time	10-15 min
CT Skills	Judge source credibility, Handle ambiguity
How to Run It	<ol style="list-style-type: none">1. Choose a factual question relevant to your teaching area.2. Search the question on Google and note the top 3-5 results and their sources.3. Ask the same question to ChatGPT and read the response.4. Compare: How many sources does each provide? How verifiable is each answer?5. Discuss: What are the trade-offs between convenience and credibility?
Source	<i>Adapted from Shroff, L. (2025). The people outsourcing their thinking to AI. The Atlantic.</i> https://www.theatlantic.com/technology/2025/12/people-outsourcing-their-thinking-ai/685093/

Activity 7: The Pricing Exercise

I picked this up from a TEDx talk by Charlie Gedeon, a university instructor and UX designer. He described a moment in class when a student priced her business at \$50 per month. When asked why, the student said, "That's what ChatGPT said." The prompt had no context, no detail about the business model, no target audience, nothing. And ChatGPT still produced a confident, specific answer.

That story is a perfect entry point for this activity. Students run a similarly vague prompt on their own devices and see what comes back. Then they evaluate: are sources cited? Is the advice specific to any real situation? What assumptions did the AI make on your behalf? After that, they rewrite the prompt with proper context and constraints and compare the two outputs. The difference is always striking.

This targets Ennis's skills of evaluating assumptions, judging source credibility, and asking clarification questions. And it makes a point that sticks: AI will always give you an answer. The question is whether that answer is worth anything without the right input. Students walk away understanding that the quality of what you get from AI depends entirely on the quality of what you ask.

Time	10-15 min
CT Skills	Evaluate assumptions, Judge source credibility, Ask clarification questions
How to Run It	<ol style="list-style-type: none"> 1. Present the scenario: a student asks ChatGPT 'How should I price my business?' with minimal context. 2. Participants run a similar vague prompt on their own devices. 3. Evaluate the output: Are sources cited? Is the advice context-specific? What assumptions does the AI make? 4. Discuss: What would a critical thinker do differently with this output? 5. Rewrite the prompt with proper context and constraints, then compare the new results.
Source	<p><i>Adapted from Gedeon, C. (2026). Can AI help us learn? TEDx Talks.</i> https://www.youtube.com/watch?v=m8WomdCLBqE</p>

Activity 8: Socratic Questioning with AI

The Socratic method has been around for centuries, and it applies naturally to AI interactions. The concept is straightforward: don't accept the first answer. Keep asking.

Students ask ChatGPT a substantive question in their field, then follow up with a series of probing questions. "What evidence supports this?" "What are you assuming here?" "What perspectives are missing?" "What would someone who disagrees say?" Each follow-up pushes the AI to reveal gaps, contradictions, or shallow reasoning that the first polished response hid. Students track how the AI's answers shift across the rounds, and that tracking is where the learning happens.

This connects to three of Ennis's skills at once: asking clarification questions, analyzing arguments, and evaluating assumptions. It also teaches a habit students can use every time they interact with a chatbot. Favero et al. (2024) explored a similar approach in their research on Socratic chatbots, showing that when AI is used as a scaffold for questioning, it can strengthen critical thinking. The key is that the student drives the conversation, not the AI.

Time	15-20 min
CT Skills	Ask clarification questions, Analyze arguments, Evaluate assumptions
How to Run It	<ol style="list-style-type: none"> 1. Ask ChatGPT a substantive question in your field of expertise. 2. Do NOT accept the first answer. Ask: 'What evidence supports this?' 3. Follow up with: 'What assumptions underlie this answer?' 4. Then ask: 'What perspectives or counterarguments are missing?' 5. Document how the AI's response evolves through each round of questioning.
Source	<i>Based on the Socratic method. See also Favero et al. (2024). Enhancing critical thinking in education by means of a Socratic chatbot. arXiv.</i>

Activity 9: AI Debate / Multiple Perspectives

This activity is commonly used in various teaching contexts, and I adapted it here based on ideas from **Nikolas Dietis** writing for **Times Higher Education**. The setup is simple. You pick a debatable topic in your discipline and prompt ChatGPT twice: once to argue for the position, once to argue against it. Each group gets one side, and their job is to fact-check the claims, spot what's missing, and add their own evidence and reasoning.

When both groups present, the full room gets to see how AI handles argumentation from both sides. It usually becomes clear pretty quickly that the AI gives you structure and fluency but lacks depth, nuance, and genuine engagement with counterarguments. The "for" essay and the "against" essay often follow the same template, and neither one wrestles with the complexity the way a real argument should.

This hits several of Ennis's skills at once: analyzing arguments, handling ambiguity, and recognizing fallacies. And it raises a question that students carry with them: if AI can argue both sides with equal confidence, what does "confidence" actually tell you about the quality of an argument?

Time	20-25 min
CT Skills	Analyze arguments, Handle ambiguity, Recognize fallacies
How to Run It	<ol style="list-style-type: none">1. Select a debatable topic relevant to your discipline.2. Prompt ChatGPT to argue FOR the position and give this output to Group A.3. Prompt ChatGPT to argue AGAINST the position and give this output to Group B.4. Each group fact-checks claims, identifies gaps, and adds its own evidence and analysis.

Source	5. Groups present their findings, and the full room discusses the quality of AI argumentation.
	<p><i>Commonly used in various teaching contexts. See also: Dietis, N. (2024, January 8). Three ways to use ChatGPT to enhance students' critical thinking in the classroom. Times Higher Education.</i></p> <p>https://www.timeshighereducation.com/campus/three-ways-use-chatgpt-enhance-students-critical-thinking-classroom</p>

Activity 10: Assignment Redesign Workshop

This one is aimed at educators, and it's the kind of activity that changes how you think about your own course design. Flower Darby wrote about a similar framework in Faculty Focus, and I adapted it here to connect with Ennis's full set of six critical thinking skills, applied not to AI output but to pedagogy itself.

You bring one assignment from your current course and break it down into its individual steps. For each step, you ask three questions: Where does AI use make sense here? Where should students work without it? And how can I build in critical thinking that AI can't shortcut? Then you redesign the assignment with transparent AI guidance at each stage, so students know exactly what's expected.

Work in small groups and give feedback on each other's redesigns. A second pair of eyes often surfaces blind spots you didn't see. Someone might point out that your "no AI" step still has a loophole, or that your "AI-allowed" step doesn't require enough critical engagement. Those conversations are where the redesign gets real. This draws on all six of Ennis's skills applied to course design, and it's one of the few activities in this guide that's specifically for the instructor.

Time	20-30 min
CT Skills	All six skills (applied to pedagogy)
How to Run It	1. Bring one assignment from your current course.

Source	<ol style="list-style-type: none">2. Break it down into its individual steps or components.3. For each step, decide: Where is AI use appropriate? Where should it be restricted? Why?4. Redesign the assignment with transparent AI guidance at each stage.5. Share redesigned assignments with peers and collect feedback. <p><i>See for instance Darby, F. (2025). 5 steps to update assignments to foster critical thinking and authentic learning in an AI age. Faculty Focus.</i></p> <p>https://www.facultyfocus.com/articles/effective-teaching-strategies/5-steps-to-update-assignments-to-foster-critical-thinking-and-authentic-learning-in-an-ai-age/</p>
---------------	--

Activity 11: Prompt Engineering for Quality

This activity makes a point that many students haven't considered: writing a good prompt is itself a critical thinking exercise. Rick Holbeck wrote about a similar approach in Faculty Focus, and I adapted it here to target two of Ennis's skills: asking clarification questions and handling ambiguity.

You start by showing a vague prompt and the AI output it produces. Then you show a refined version of the same prompt, one with context, constraints, and a clear format, and compare what comes back. The difference is always obvious. The vague prompt produces generic filler. The refined prompt produces something you can actually work with.

Then participants write their own prompts for a real teaching scenario they face, test them live on ChatGPT, and compare results with the person next to them. The conversation naturally turns to why specificity matters, and that connects to Ennis's framework. If you can't ask a clear question, you can't expect a useful answer, from AI or from anyone. Students walk away from this one with a practical skill they can use immediately, and a deeper understanding of why precision in communication matters.

Time	15-20 min
CT Skills	Ask clarification questions, Handle ambiguity
How to Run It	<ol style="list-style-type: none"> 1. Show a vague prompt and its AI output side by side. 2. Show a refined prompt (with context, constraints, and desired format) and its improved output. 3. Participants write their own improved prompts for a teaching scenario they face. 4. Test prompts live on ChatGPT and compare the quality of results. 5. Discuss: How does the process of writing a better prompt itself require critical thinking?
Source	<p>See for instance Holbeck, R. (2025, June 27). <i>Helping students develop AI prompting skills for critical thinking</i>. <i>Faculty Focus</i>.</p> <p>https://www.facultyfocus.com/articles/teaching-with-technology-articles/helping-students-develop-ai-prompting-skills-for-critical-thinking/</p>

Activity 12: Funhouse Mirror

This one comes from Roe, Furze, and Perkins (2024), who use four metaphors to describe how AI handles information, and the funhouse mirror is one of them. The idea is simple: pick a topic you know deeply and ask ChatGPT to explain it. Then read the output carefully.

You'll recognize the shape of your subject, but the details will be off. Key nuances get flattened, important differences between concepts disappear, and confident-sounding sentences paper over gaps that any expert would catch immediately. That's the funhouse mirror effect: the reflection looks close enough to pass, but it's distorted in ways that matter.

This connects to Ennis's skills of handling ambiguity, evaluating assumptions, and recognizing fallacies. And it raises an important follow-up question: if AI can mislead you on a topic you've spent years studying, what happens when students encounter topics they're seeing for the first time? That question tends to land hard, and it's the kind of realization that changes how students approach AI output going forward.

Time	10-15 min
CT Skills	Handle ambiguity, Evaluate assumptions, Recognize fallacies
How to Run It	<ol style="list-style-type: none"> 1. Choose a topic you are genuinely expert in. 2. Ask ChatGPT to explain or summarize that topic in detail. 3. Read the output carefully: What is accurate? What is distorted or oversimplified? What is missing entirely? 4. Write down specific distortions and explain why they matter in your field. 5. Discuss: If AI can mislead you on your own topic, what happens when students encounter unfamiliar topics?
Source	<i>Roe, J., Furze, L., & Perkins, M. (2024). Funhouse mirror or echo chamber? A methodological approach to teaching critical AI literacy through metaphors. arXiv.</i>

Activity 13: Critical Evaluation Checklist (The Reveal)

This one has a twist, and it's one of the most effective activities in this guide for surfacing students' own assumptions. I adapted it from a teaching resource published by the University at Albany.

You give participants a text and tell them it's a student submission or a draft report. They evaluate it using a structured checklist: Is it accurate? Is the logic sound? Are sources cited and credible? Is there bias? Is anything missing? They score it, take notes, and form an opinion. Then you reveal that the entire text was generated by ChatGPT.

The conversation that follows gets at something important: do we evaluate content differently depending on who we think produced it? And if so, what does that say about our own assumptions? Most people find that they were more generous in their initial assessment because they assumed a human wrote it. That gap between how we evaluate human work and how we evaluate AI work is exactly the kind of blind spot that critical thinking training needs to address. This activity touches all six of Ennis's skills, but the real power is in the debrief.

Time	15-20 min
CT Skills	All six critical thinking skills
How to Run It	<ol style="list-style-type: none">1. Distribute an AI-generated text presented as a 'student submission' or 'draft report.'2. Participants evaluate using a checklist: accuracy, logic, sources, bias, completeness.3. Collect initial evaluations and note scores.4. Reveal: this text was generated entirely by ChatGPT.5. Discuss: Did knowing the source was AI change your assessment? Why or why not?
Source	<i>University at Albany. (n.d.). Help students be critical of AI.</i> https://www.albany.edu/teaching-and-learning/teaching-resources/help-students-be-critical-ai

Conclusion

I Hope these activities give you something practical to work with the next time you walk into your classroom. They're not meant to be followed rigidly. Adapt them, combine them, adjust the timing, swap the AI tools, change the prompts to fit your discipline. The framework is flexible on purpose.

But I want to leave you with one thought that runs through everything in this guide. The challenge with generative AI in education isn't the technology itself. It's the default mode most students fall into when they use it: prompt, receive, accept, move on. That cycle is where critical thinking goes quiet. Every activity in this collection is designed to break that cycle, to insert a moment of friction where students have to stop, question, evaluate, and respond with their own reasoning before they move forward.

That's not anti-AI. It's pro-thinking and there's a big difference.

Teaching in the age of AI is hard. It asks us to rethink assignments we've used for years, to learn tools that didn't exist eighteen months ago, and to hold two things at once: openness to what AI can do and healthy skepticism about what it can't. If you've read this far, you're already doing that work. Keep going.

About the Author

Med Kharbach, PhD, is an educator, researcher, and the editor of *Educators Technology* (educatorstechnology.com). A former K-12 teacher with nearly two decades of teaching experience across K-12 and higher education, Med currently serves as an Instructor at Mount Saint Vincent University, where he teaches Critical Theory and Education at the graduate level. 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* and *The BEARA Framework for Pedagogical Integration* with Dr. Jonathan Woodworth. His work focuses on AI literacy, assessment in the age of AI, and practical classroom applications of educational technology. You can read more about Med and his research at medkharbach.com.

References

Abdulhai, M., White, I., Wan, Y., Qureshi, I., Leibo, J., Kleiman-Weiner, M., & Jaques, N. (2026). How LLMs distort our written language. arXiv preprint arXiv:2603.18161. <https://arxiv.org/abs/2603.18161>

- Barcaui, A. (2025). ChatGPT as a cognitive crutch: Evidence from a randomized controlled trial on knowledge retention. *Social Sciences & Humanities Open*, 12, 102287. <https://doi.org/10.1016/j.ssaho.2025.102287>
- Bastani, H., Bastani, O., Sungu, A., Geb, H., Kabakcı, Ö., & Marimane, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122> //
- Darby, F. (2025, November 5). 5 steps to update assignments to foster critical thinking and authentic learning in an AI age. *Faculty Focus*. <https://www.facultyfocus.com/articles/effective-teaching-strategies/5-steps-to-update-assignments-to-foster-critical-thinking-and-authentic-learning-in-an-ai-age/>
- Dietis, N. (2024, January 8). Three ways to use ChatGPT to enhance students' critical thinking in the classroom. *Times Higher Education*. <https://www.timeshighereducation.com/campus/three-ways-use-chatgpt-enhance-students-critical-thinking-classroom>
- Ennis, R. H. (2015). Critical thinking: A streamlined conception. In M. Davies & R. Barnett (Eds.), *The Palgrave handbook of critical thinking in higher education* (pp. 31-47). Palgrave Macmillan.
- Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, 56(2), 489–530. <https://doi.org/10.1111/bjet.13544>
- Gedeon, C. (2026). Can AI help us learn? [Video]. *TEDx Talks*. <https://www.youtube.com/watch?v=m8WomdCLBqE>
- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), Article 6. <https://doi.org/10.3390/soc15010006>
- Favero, L., Pérez-Ortiz, J. A., Käser, T., & Oliver, N. (2024, October). Enhancing critical thinking in education by means of a Socratic chatbot. In *International workshop on AI in education and educational research* (pp. 17-32). Cham: Springer Nature Switzerland. <https://arxiv.org/html/2409.05511v1>
- Holbeck, R. (2025, June 27). Helping students develop AI prompting skills for critical thinking. *Faculty Focus*. <https://www.facultyfocus.com/articles/teaching-with-technology/articles/helping-students-develop-ai-prompting-skills-for-critical-thinking/>
- Kingston, M. (2023). Hallucination detective [Lesson plan]. *AI for Education*. <https://www.aiforeducation.io/curriculum/lesson-3>
- Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing tasks. MIT Media Lab. <https://www.media.mit.edu/publications/your-brain-on-chatgpt/>
- McMahon, K. (2025). Critical thinking in the age of AI [Online course]. *Maven Analytics*. <https://mavenanalytics.io/course/critical-thinking-in-the-age-of-ai>
- Means, T. (2025, December 22). The push-back protocol: Teaching students to challenge AI, not accept it. *The Collaboration Chronicle: Human+AI in Education*. <https://tawnyameans.substack.com/p/the-push-back-protocol-teaching-students>
- Pardo-Guerra, J. P. (n.d.). Using AI text as prompts for critical analysis. In *TextGenEd collection*. *WAC Clearinghouse*.

- <https://wacclearinghouse.org/repository/collections/textgened/rhetorical-engagements/using-ai-text-as-prompts-for-critical-analysis/>
- Roe, J., Furze, L., & Perkins, M. (2024). Funhouse Mirror or Echo Chamber? A Methodological Approach to Teaching Critical AI Literacy Through Metaphors. *arXiv*.
<https://arxiv.org/abs/2411.14730>
- Shroff, L. (2025, December 1). The people outsourcing their thinking to AI. *The Atlantic*.
<https://www.theatlantic.com/technology/2025/12/people-outsourcing-their-thinking-ai/685093/>
- Sourati, Z., Ziabari, A. S., & Dehghani, M. (2026). The homogenizing effect of large language models on human expression and thought. *Trends in Cognitive Sciences*. Advance online publication. <https://doi.org/10.1016/j.tics.2026.01.003>
- University at Albany. (n.d.). *Help students be critical of AI*. <https://www.albany.edu/teaching-and-learning/teaching-resources/help-students-be-critical-ai>