

Critical Thinking Activities for the Age of AI

Hands-On Exercises for Faculty and Students

Dr. Med Kharbach | NENA ICHRIE 2026

This handout presents a collection of hands-on activities designed to help faculty integrate critical thinking into AI-assisted teaching. Each activity is grounded in Robert Ennis's (2015) Critical Thinking Skills framework and targets specific skills such as analyzing arguments, judging source credibility, handling ambiguity, evaluating assumptions, recognizing fallacies, and asking clarification questions.

These activities can be easily adapted to various AI-assisted teaching contexts, disciplines, and student levels. Use them as standalone exercises, combine them into workshop sequences, or embed them into your existing course design. The goal is to equip students (and ourselves) with the thinking skills needed to engage with AI critically.

The Six CT Skills (Ennis, 2015, pp. 32-33):

- Analyze arguments
- Judge source credibility
- Handle ambiguity
- Evaluate assumptions
- Recognize fallacies
- Ask clarification questions

Activity 1: Push-Back Protocol

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

question relevant to their discipline, then go through multiple rounds of structured push-back. They demand evidence for claims, question what's being assumed, look for alternative perspectives, and stress-test the reasoning. Each round maps to a specific skill on the rubric, so what starts as a simple chatbot exchange becomes an active critical thinking exercise. Faculty can use this with students at any level and in any discipline.

Activity 1: Push-Back Protocol	
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. Push back: ask for evidence, sources, and reasoning behind each claim. 4. Challenge assumptions: what is being taken for granted in the response? 5. Seek alternatives: ask AI if there are opposing views or exceptions.
Source	<p>Means, T. (2025). <i>The Push-Back Protocol: Teaching Students to Challenge AI, Not Accept It</i>. Substack.</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, and I adapted it here to focus on Ennis's skill of judging source credibility. Participants receive an AI-generated text with errors baked in: fake citations, wrong dates, made-up statistics, and confident-sounding claims with no basis. Their job is to play detective, go through the text line by line, flag what looks suspicious, and verify claims against reliable databases. You can open with the [Steven Schwartz case](#) (the lawyer who submitted a ChatGPT brief full of fabricated citations) to set the stakes, but the real learning comes in the debrief: why does AI hallucinate, and what habits help us catch it?

Activity 2: Hallucination Detective

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. Participants 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]</i>. <i>AI for Education</i>. https://www.aiforeducation.io/curriculum/lesson-3</p>

Activity 3: Human vs. AI Blind Comparison

This one is simple and effective. 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. Participants read both and try to figure out which is which. The conversation that follows is where the value is. People start noticing patterns: the AI text tends to sound polished but vague, sources are missing or hard to trace, and the reasoning often stays surface-level. This connects directly to several of Ennis's skills, especially judging source credibility, analyzing arguments, and handling ambiguity. You can run this in any discipline with minimal preparation.

Activity 3: Human vs. AI Blind Comparison

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 this handout set, which I built based on Ennis's (2015) six critical thinking skills and inspired by Maven Analytics' course on Critical Thinking in the Age of AI. Participants receive a piece of AI-generated text and score it across all six standards: argument analysis, source credibility, ambiguity, assumptions, fallacies, and clarification.

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. You can use this with any AI-generated content in any discipline.

Activity 4: AI Output Audit with Rubric	
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	<p>McMahon, K. (2025). <i>Critical thinking in the age of AI</i> [Online course]. Maven Analytics. https://mavenanalytics.io/course/critical-thinking-in-the-age-of-ai</p>

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. I adapted it to connect with Ennis's skills of analyzing arguments, evaluating assumptions, and recognizing fallacies. Participants receive an AI-generated essay or report and go through it with a pen (or digitally with comments using tools 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 forces precision. You can't just say "this feels off." You have to name the problem and propose a fix.

Activity 5: Text Annotation Exercise

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

This one takes two minutes to set up and always sparks a good conversation. Participants 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. Participants start asking: which answer would a student trust more? And 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 from earlier in the workshop: just because something reads well doesn't mean it's accurate.

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

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	<p><i>Adapted from Shroff, L. (2025). The people outsourcing their thinking to AI. The Atlantic.</i></p> <p>https://www.theatlantic.com/technology/2025/12/people-outsourcing-their-thinking-ai/685093/.</p>

Activity 7: The Pricing Exercise

I picked this up from a TED talk where Charlie Gedeon, a university instructor and UX designer, described a moment in class when a student priced her business at \$50 per month. When asked why, she 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. I turned this into an activity that targets Ennis's skills of evaluating assumptions, judging source credibility, and asking clarification questions.

Participants 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 compare the two outputs. The difference is always striking.

Activity 7: The Pricing Exercise

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></p> <p>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. Participants 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. Participants track how the AI's answers shift across the rounds. 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.

Activity 8: Socratic Questioning with 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. Instead, 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	Based on the Socratic method. See also Favero et al. (2024). <i>Enhancing Critical Thinking in Education by Means of a Socratic Chatbot</i> . arXiv. https://arxiv.org/html/2409.05511v1

Activity 9: AI Debate / Multiple Perspectives

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 that the AI gives you structure and fluency but lacks depth, nuance, and genuine engagement with counterarguments. This one hits several of Ennis's skills at once: analyzing arguments, handling ambiguity, and recognizing fallacies.

Activity 9: AI Debate / Multiple Perspectives

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

Activity 10: Assignment Redesign Workshop

Bring one assignment from your current course and break it down into its individual steps. For each step, ask: 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? 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. This activity draws on all six of Ennis's skills applied to course design.

Activity 10: Assignment Redesign Workshop	
Time	20-30 min
CT Skills	All six skills (applied to pedagogy)
How to Run It	<ol style="list-style-type: none"> 1. Bring one assignment from your current course. 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.
Source	<p>See for instance Darby, F. (2025). <i>5 steps to update assignments to foster critical thinking and authentic learning in an AI age</i>. <i>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 often gets overlooked: writing a good prompt is itself a critical thinking exercise. Start by showing a vague prompt and the AI output it produces. Then 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. 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 two of Ennis's skills: asking clarification questions and handling ambiguity. If you can't ask a clear question, you can't expect a useful answer, from AI or from anyone.

Activity 11: Prompt Engineering for Quality

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>. 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 distinctions 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?

Activity 12: Funhouse Mirror	
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	<p>Roe, J., Furze, L., & Perkins, M. (2024). <i>Funhouse Mirror or Echo Chamber? A Methodological Approach to Teaching Critical AI Literacy Through Metaphors</i>. <i>arXiv</i>. https://arxiv.org/abs/2411.14730</p>

Activity 13: Critical Evaluation Checklist

This one has a twist. 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? This activity touches all six of Ennis's skills, but the real power is in the debrief.

Activity 17: Critical Evaluation Checklist (The Reveal)	
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	University at Albany. (n.d.). <i>Help students be critical of AI</i> . https://www.albany.edu/teaching-and-learning/teaching-resources/help-students-be-critical-ai

References

- 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.
- Gedeon, C. (2026). Can AI help us learn? [Video]. *TEDx Talks*.
<https://www.youtube.com/watch?v=m8WomdCLBqE>
- 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>
- 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/>

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>