AI Fact-Checking Activities for Students

Classroom Activities with Printable Worksheets

Med Kharbach, PhD
medkharbach.com

Introduction

In my book *Critical Thinking with AI: 81 Practical Activities for Teachers*, I argued that evaluation deserves a different place in how we think about learning. In the revised version of Bloom's taxonomy (Anderson & Krathwohl, 2001), Evaluate sits just below Create, near the top of the hierarchy. For a long time that order made sense. It was hard to produce an original piece of writing, and it seemed natural to treat creation as the most demanding thing a learner could do.

Generative AI has changed that. For the first time in human history, we have a technology that produces long stretches of text that resemble human language to a remarkable degree. Anyone with a device can now generate an essay, a report, or a story in seconds. On the surface at least, creation has become easy. What's become harder, and more important, is judging what gets created. That's why I believe Evaluate now belongs at the same level as Create.

Assessment researchers have a name for this capability: evaluative judgement, the ability to make decisions about the quality of one's own work and the work of others (Tai et al., 2018). Bearman and colleagues (2024) argue that developing it has become even more urgent now that students work alongside generative AI. When a tool can produce a fluent answer on demand, the student's job shifts to deciding whether that answer is accurate, complete, and worth using.

Evaluation also brings several literacies together. When students check an AI answer, they draw on information literacy to trace a claim to its source, and on media literacy to ask who produced a piece of content and why. Digital literacy comes in as they search, compare, and read laterally across websites. AI literacy helps them understand why a chatbot can sound confident while being wrong. Evaluation also calls on a range of thinking skills: analyzing a passage into separate claims, searching for evidence, verifying details, weighing sources against each other, and synthesizing what they find into a judgement they can defend.

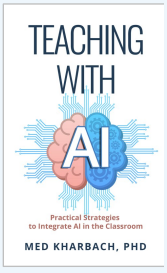
This guide turns those ideas into classroom practice. I compiled it with the help of Claude. It contains ten activities for grades 6 to 12, several of which also work in first-year university courses. Each one includes teacher steps, a practice passage written for the lesson, a list of claims to pay attention to, and a printable student worksheet. The activities build on each other. They start with breaking a paragraph into separate claims, move through checking citations, sources, numbers, and quotations, and end with examining how the wording of a question shapes the answer.

One request before you start: check the claims in each activity yourself before class, using current and reliable sources. Some facts change over time, and the debrief works best when you know where the evidence stands.

I hope this short guide helps your students build the habits they'll need to question what AI tells them, and the confidence to trust their own judgement when they do.

More from Med Kharbach

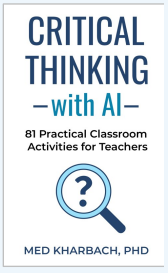
Two books for teachers who want to go further with AI in the classroom



Teaching with AI

Practical strategies to integrate AI in the classroom, from building your own AI literacy to aligning AI use with sound pedagogy and professional judgement.

Get your copy



Critical Thinking with AI

81 classroom activities that ask students to question claims, weigh evidence, and explain their reasoning, from Grade 3 through university.

Get your copy

Activity 1: Claim Hunt

Grade band: 6 to 12 (adaptable for upper elementary)
Time: 40 to 50 minutes (about 70 minutes if students also verify claims)

Purpose

Fact-checking becomes more manageable when students break a paragraph into individual claims. This activity helps them identify statements that can be checked against evidence, distinguish them from opinions, and decide where checking should begin.

The learning outcome is simple: students identify separate claims and explain what evidence they would need to check them. Verifying those claims comes later in the guide, though you can extend this lesson to include it.

Materials

- The teacher-prepared practice passage below, or a short AI-generated answer on a topic your class is studying
- Claim Hunt worksheet (one per student or pair)
- Pencils, plus highlighters if you want students to mark claim types
- Optional, if students will verify claims in this lesson: device access, about 20 extra minutes of research time, and the evidence column on the worksheet to record what they find

Key Terms for Students

Students label each claim in two ways. First, its **type**:

- **Factual claim:** a statement about the world that evidence can confirm or challenge. A factual claim can turn out to be accurate, inaccurate, misleading, or unsupported. Calling it a factual claim says nothing yet about whether it is true.
- **Opinion or value judgement:** a statement of preference or evaluation, such as "best" or "most fascinating."
- **Prediction:** a statement about what will happen.
- **Mixed or unclear:** a statement that blends types or is too vague to classify.

Second, whether it can be **checked**:

- **Checkable now:** evidence exists that we could look for today.
- **Needs clarification:** the wording has to be pinned down before anyone can check it.
- **Opinion or preference:** evidence can inform it but cannot settle it.
- **Future outcome:** the event hasn't happened, though we can still ask what evidence supports the prediction.

Teacher Steps

1. **Open with a question.** Ask students: "If an AI gives you a paragraph about a topic, how much of it would you check?".
2. **Read the passage aloud.** Hand out the practice passage and read it together once without stopping. Tell students it was written for this lesson and contains a mix of accurate, inaccurate, unsupported, and opinion statements.
3. **Model splitting a sentence.** Take the waggle dance sentence and mark it with slashes as you think aloud: "Honeybees use a waggle dance to share where flowers are / Karl von Frisch helped establish what the dance means / he won the Nobel Prize in 1973. That's three claims in one sentence, and each one would need its own check." Point out that qualifying words such as "only," "up to," and "never" must stay in the claim when students paraphrase, because they change what needs checking.
4. **Introduce the labels.** Walk through the claim types and checkability options above, using one example of each from the passage.
5. **Students hunt in pairs.** Pairs complete Parts A and B of the worksheet, listing every claim they can find and noting what evidence they would look for.
6. **Discuss as a class.** Ask: Which claims were easiest to identify? Which would require clarification before you could check them? What evidence would you look for?
7. **Close with corrections.** Go through the claims below with the class, using the sources you checked, so that no false or unsupported claim is left hanging. Students should leave knowing which statements were wrong and why.
8. **Individual reflection.** Students complete Parts C and D on their own.

Teacher-Prepared Practice Passage

This passage was written for teaching. It contains accurate, inaccurate, unsupported, and opinion statements.

> *Western honeybees (Apis mellifera) are the only insects that make honey. A single western honeybee colony can hold up to 60,000 bees during the summer months. Honeybees tell each other where to find flowers through a "waggle dance," whose meaning Karl von Frisch helped establish through his research, work that earned him the Nobel Prize in 1973. Honey never spoils, and archaeologists have found edible honey in Egyptian tombs more than 3,000 years old. Western honeybees are native to North America and have lived there for millions of years. Scientists warn that without bees, humans would die out within four years. Bees are the most fascinating creatures on Earth.*

Claims to Pay Attention To

These are the claims in the practice material that deserve the closest attention. Check them yourself before class, using current and reliable sources, so you can guide the debrief and prepare a source pack if your students need one.

- **"The only insects that make honey":** The word "only" is what makes this checkable. Are any other insects known to make honey?
- **"Up to 60,000 bees":** Check colony size, and watch whether students keep "up to" when they paraphrase.
- **The waggle dance sentence:** Three claims in one sentence: the dance itself, Karl von Frisch's role in explaining it, and the 1973 Nobel Prize.
- **"Honey never spoils":** The word "never" hides conditions. What does honey's resistance to spoilage depend on?
- **Edible honey in Egyptian tombs:** Widely repeated, but can it be traced to an original archaeological source?
- **"Native to North America":** Where did western honeybees in North America come from, and when?
- **The four-year warning:** Two claims: that scientists issued this warning, and that humans would die out within four years. Who said it, and what evidence supports it?
- **"The most fascinating creatures on Earth":** An opinion or value judgement, not a factual claim.

Extension

Have students generate their own AI answer on a current unit topic and swap with another pair, who then run the Claim Hunt on it. For younger students, cut the passage to four sentences and use only the claim type column.

Student Worksheet: Claim Hunt

Name: _____ Date: _____

Part A: Find the Claims

Underline each separate claim in the passage. A sentence may contain more than one. Use slashes (/) to show where you would separate them.

When you rewrite a claim in your own words, keep qualifying words such as **"only," "up to," "never,"** and **"always."** They change what needs checking.

Claim types: Factual claim / Opinion or value judgement / Prediction / Mixed or unclear
Can it be checked? Checkable now / Needs clarification / Opinion or preference / Future outcome

Part B: List the Claims

#	The claim in my own words	Claim type	Can it be checked?	What evidence would I look for?
1				
2				
3				
4				
5				
6				
7				
8				

Found more claims? Continue on the back of this page.

Part C: Plan Your Checks

A claim I would check first, and why:

A claim that needs clearer wording before it can be checked:

A statement that expresses an opinion:

Evidence I would need before accepting one of the claims:

Part D: Reflect

1. Did any sentence contain more than one claim? Which one, and how did you split it?

2. What clue words helped you spot an opinion (for example, "best," "most," "amazing")?

3. Which claim feels familiar? Could that familiarity affect your judgement?

4. If you could ask the AI one follow-up question, what would it be?

Activity 2: Citation Check

Grade band: 8 to 12 (works well in first-year university too)
Time: 45 to 60 minutes

Purpose

AI tools often attach references to their answers, and a plausible-looking citation can make an answer seem trustworthy. A reference can point to a study nobody can locate, a real study with the wrong details, or a real study that says something different from what the AI claims. This activity teaches students to open every source and run three checks: can I locate it, do the details match, and does it support the claim?

Materials

- The teacher-prepared practice answer below
- Citation Check worksheet (one per pair)
- Devices with internet access (Google Scholar, library databases, publisher sites, or a regular search engine)

Teacher Steps

1. **Introduce the citation problem.** Explain that AI tools sometimes produce references that look completely real, with authors, journal names, page numbers, and even DOIs, yet contain errors or point to sources nobody can locate. Ask: "How would you know whether a source is real and says what the AI claims?" Collect ideas on the board.
2. **Introduce the three checks.** Write them up and keep them visible for the whole lesson:

- **Exists:** Can I locate this source outside the AI?
- **Matches:** Are the authors, year, title, journal, and publication status correct?
- **Supports:** Does the source say what the AI claims it says?

3. **Teach a search sequence.** Students search in this order: (1) the exact title in quotation marks, (2) author names plus keywords, (3) publisher, journal, or repository records such as Google Scholar, PubMed, or arXiv. Set a limit of about five minutes per reference so no pair spends the lesson hunting for a source that can't be found.

4. **Model the Sparrow, Liu, and Wegner reference.** Search the title, locate the publication record, and compare its details with the reference provided. Think aloud as you compare each detail: authors, year, title, journal, volume, and pages.
5. **Pairs check the rest.** Pairs work through the remaining references, recording where they searched, what they found, and the evidence behind each judgement.
6. **Compare verdicts.** Go reference by reference as a class. Pairs will disagree on some, and those disagreements make for the best discussion.
7. **Debrief with corrections.** Go through each reference with the class, using the publication records you checked, so students leave knowing what held up and what didn't.
8. **Reflect.** Students complete Part C individually.

Teacher-Prepared Practice Answer

This AI-generated answer was written for teaching. It contains deliberately altered and invented references.

> *Research shows that technology has a strong effect on how students learn and remember. Sparrow, Liu and Wegner (2015) found that people are less likely to remember information when they expect to be able to look it up later. Mueller and Oppenheimer (2014) showed that students who typed their notes on laptops performed better on conceptual questions than students who wrote by hand. A national survey by Thompson and Alvarez (2023) found that 72% of teenagers lose focus within four minutes of receiving a phone notification. More recently, a peer-reviewed MIT study by Kosmyna et al. (2025) found that students who used ChatGPT to write essays showed weaker brain connectivity than students who wrote without help.*
>
> ***References***
>
> *Kosmyna, N., et al. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. MIT Media Lab.*
>
> *Mueller, P. A., & Oppenheimer, D. M. (2014). The pen is mightier than the keyboard: Advantages of longhand over laptop note taking. Psychological Science, 25(6), 1159-1168.*
>
> *Sparrow, B., Liu, J., & Wegner, D. M. (2015). Google effects on memory: Cognitive consequences of having information at our fingertips. Nature, 521, 776-778.*

> *Thompson, R., & Alvarez, M. (2023). Smartphones and adolescent attention: A national survey. Journal of Youth Digital Studies, 14(2), 88-104.*

Claims to Pay Attention To

These are the claims in the practice material that deserve the closest attention. Check them yourself before class, using current and reliable sources, so you can guide the debrief and prepare a source pack if your students need one.

- **Sparrow, Liu, and Wegner:** Compare the year, journal, volume, and pages with the publication record students locate.
- **Mueller and Oppenheimer:** The citation details may look right. Check whether the finding matches what the study actually reported, and note that students are checking one study, not settling whether handwriting is always better.
- **Thompson and Alvarez:** Check whether this exists. Students should report whether they could locate it and document their searches. An unsuccessful search alone does not prove a source is fabricated.
- **Kosmyna et al.:** Check the publication status the answer claims ("peer-reviewed") against the version cited, and how far the finding can be interpreted. Recheck the status whenever you update the guide.

A line to leave students with: a reference list is a set of claims too, and each one can be checked.

Extension

Ask an AI tool to answer a question from your current unit and include five references. Students run all three checks on each one and explain which checks each reference passed, with evidence. As a follow-up, they can ask the AI directly, "Are these sources real?" Make one point explicit: the AI's reassurance is another claim to check. It is not independent verification.

Student Worksheet: Citation Check

Name: _____ Date: _____

Part A: The Three Checks and Search Sequence

For every reference, ask:

1. **Exists:** Can I locate it outside the AI?
2. **Matches:** Are the authors, year, title, journal, and publication status correct?
3. **Supports:** Does the source say what the AI claims?

Search in this order: exact title in quotation marks, then author names plus keywords, then publisher, journal, or repository records. Spend no more than about five minutes on each reference.

Part B: Check Each Reference

Reference 1

Author(s) and year	
Where I searched	
Located?	☐ Located ☐ Not located ☐ Uncertain
Details correct? What is wrong?	
Evidence: link and relevant passage or section	
Does it support the AI's claim?	☐ Supports ☐ Partly supports ☐ Contradicts ☐ Cannot determine
My verdict (tick all that apply)	☐ Located and accurate ☐ Incorrect details ☐ Not located ☐ Findings misrepresented ☐ Publication status misrepresented

Reference 2

Author(s) and year	
Where I searched	
Located?	☐ Located ☐ Not located ☐ Uncertain
Details correct? What is wrong?	
Evidence: link and relevant passage or section	
Does it support the AI's claim?	☐ Supports ☐ Partly supports ☐ Contradicts ☐ Cannot determine
My verdict (tick all that apply)	☐ Located and accurate ☐ Incorrect details ☐ Not located ☐ Findings misrepresented ☐ Publication status misrepresented

Reference 3

Author(s) and year	
Where I searched	
Located?	☐ Located ☐ Not located ☐ Uncertain
Details correct? What is wrong?	
Evidence: link and relevant passage or section	
Does it support the AI's claim?	☐ Supports ☐ Partly supports ☐ Contradicts ☐ Cannot determine
My verdict (tick all that apply)	☐ Located and accurate ☐ Incorrect details ☐ Not located ☐ Findings misrepresented ☐ Publication status misrepresented

Reference 4

Author(s) and year	
Where I searched	
Located?	☐ Located ☐ Not located ☐ Uncertain
Details correct? What is wrong?	
Evidence: link and relevant passage or section	
Does it support the AI's claim?	☐ Supports ☐ Partly supports ☐ Contradicts ☐ Cannot determine
My verdict (tick all that apply)	☐ Located and accurate ☐ Incorrect details ☐ Not located ☐ Findings misrepresented ☐ Publication status misrepresented

Part C: Reflect

1. Which reference was hardest to check, and why?

2. Which features made a reference look trustworthy before you checked it?

3. How could a real source be used to support a claim it does not actually make?

4. Write one rule for yourself about using AI-provided references in your own work.

Activity 3: Lateral Reading Relay

Grade band: 7 to 12 (works in first-year university too)
Time: 50 to 60 minutes

Purpose

When students check an AI answer, they often reread the answer itself or ask the AI whether it's right. Neither tells them much. Professional fact-checkers leave the page and look at what other sources say about the claim and about the people behind it. This activity turns that habit, called lateral reading, into a team relay. Everyone works on every move, and a different student leads and records each round.

Background for Teachers

The idea comes from research by Sam Wineburg and Sarah McGrew of the Stanford History Education Group, first released as a report in 2017 and later published in Teachers College Record. They observed 10 professional fact-checkers, 10 historians, and 25 students as they evaluated online information (Stanford Report). The fact-checkers were the fastest and most accurate.

They scanned a site briefly, then opened new browser tabs to seek context and perspective from other sites, a practice the researchers called lateral reading. Historians and students tended to read vertically, staying on the original site, and were often taken in by a professional-looking name and logo, scholarly references, or a nonprofit web address. Fact-checkers also showed click restraint, looking over search results carefully before choosing one. The sample was small and deliberately made up of people expected to be skilled, so the study shows what expert strategies look like more than how typical users behave.

The moves students practise come from Mike Caulfield's SIFT method: **Stop, Investigate the source, Find better coverage, Trace claims to their original context** (Caulfield's original SIFT explanation). In this activity, students work through the moves in order. That sequence is a teaching scaffold. Caulfield's account allows different routes depending on the task, and in independent checking students can revisit moves or start with whichever one addresses their uncertainty.

AI answers add a twist. There's often no website to investigate, only the sources the AI provides, or none at all. And the signals that fooled the vertical readers in the Stanford study, such as confident wording and scholarly-looking references, are exactly what AI answers

tend to supply. So students apply lateral reading to the claims and to any sources mentioned, and look for coverage outside the AI. Asking the AI to confirm its own answer doesn't count as leaving the page.

Materials

- The teacher-prepared practice answer below
- Lateral Reading Relay worksheet (one per team)
- Devices with internet access (one per student if possible)
- A timer
- **For Grade 7 and 8:** a fallback source pack you assemble from reliable sources on the claims listed under Claims to Pay Attention To. Students attempt the searches first, then use the pack if access or technical language becomes a barrier.

Teacher Steps

Suggested timing: introduction and modelling, 10 minutes; Stop, 3 minutes; Investigate, Find, and Trace, 7 minutes each; team verdicts, 8 minutes; debrief and individual reflection, 10 minutes.

1. **Start from students' own habits.** Display the practice answer and ask: "How would you check this?" Record their responses on the board without judging them. Then name two habits. Staying within the original answer is vertical reading. Searching elsewhere begins lateral reading, but students still need to choose results carefully.
2. **Introduce lateral reading.** Share the Stanford finding in two or three sentences. Point out that the fact-checkers read less of the original page and learned more by checking what other sites said about it. Add a teaching point of your own: the most useful sources are independent of each other, not several copies of the same press release.
3. **Teach and model the four SIFT moves.** Write them on the board with their guiding questions (see Part A of the worksheet). Model one quick lateral search on a claim from outside the passage, showing click restraint: look over the results before choosing one. Explain that the fixed order is a scaffold for today.
4. **Form teams of three or four.** Everyone searches and contributes in every round. Each round, a different student leads the discussion and records in the relay log.

5. **Stop together.** Teams spend three minutes on the Stop questions before anyone searches.
6. **Run the relay.** Teams work through Investigate, Find better coverage, and Trace, about seven minutes each, rotating the lead and recorder at each move. Each new leader reads the previous notes before the team starts.
7. **Reach team verdicts.** Teams rate each claim in Part C, citing evidence from the relay log. A justified "cannot determine" is a good answer when the evidence doesn't settle the question.
8. **Debrief with corrections and reflect.** Go through the claims below, using the sources you checked, so every claim is resolved before students leave. Make one point explicit: knowing who funded a study is context for judging it, and it doesn't on its own make the findings wrong. Students then complete Part D individually.

Teacher-Prepared Practice Answer

This AI-generated answer was written for teaching. It contains accurate, overstated, underspecified, and inaccurate statements.

> *Chocolate milk is the best recovery drink for athletes, according to science. A 2006 study from Indiana University found that cyclists who drank chocolate milk after a hard workout lasted longer in their next workout than those who drank a sports drink. Since then, dozens of studies have confirmed these results, and experts now agree that chocolate milk outperforms every other recovery drink. The 2006 study was independent research with no industry involvement.*

Claims to Pay Attention To

These are the claims in the practice material that deserve the closest attention. Check them yourself before class, using current and reliable sources, so you can guide the debrief and prepare a source pack if your students need one.

- **"The best recovery drink, according to science":** Best for whom, after what exercise, compared with which alternatives, measured by what outcome?
- **The 2006 Indiana University study:** Which drinks did it compare, what did it find for each comparison, and who took part? Does "a sports drink" describe the comparison accurately?
- **"Dozens of studies have confirmed these results":** How many studies exist, and how recent is any review students rely on?

- **"Experts now agree":** What do reviews say about the strength and limits of the evidence?
- **"Outperforms every other recovery drink":** Does the evidence support superiority over all alternatives?
- **"Independent research with no industry involvement":** What does the study's own funding acknowledgment say? Funding is context: it doesn't, on its own, show whether a sponsor influenced the design, analysis, or reporting.

Extension

Give teams a fresh AI answer on a current unit topic and run the relay again, this time letting teams choose which move to start with based on what they are least sure about. Afterward, ask them to explain why they chose that route. For older students, add a fifth question to Part D: "Which sources you found were actually independent of each other, and how could you tell?"

Student Worksheet: Lateral Reading Relay

Team members: _____ Date: _____

Part A: The Four Moves

1. **Stop:** What exactly is being claimed? What do we already know about the source? Which claim needs checking first, and how much checking does this task require?
2. **Investigate the source:** Who is behind the source the AI mentions? What do others say about them?
3. **Find better coverage:** What do other sources say about the same claim? Are those sources independent of each other? Look over the search results before you click.
4. **Trace:** Where did the claim originally come from, and what did that source actually say?

Relay rules: Everyone contributes in every round. Rotate the leader and recorder at each move. Leave the AI answer: asking the AI to check itself doesn't count.

Part B: Relay Log

Move	Leader and recorder	What we did or searched	What we found (with link)	What we're passing on
1. Stop				
2. Investigate the source				
3. Find better coverage				
4. Trace				

Part C: Team Verdict

Rating options: Supported / Partly supported / Overstated / Underspecified / Contradicted / Cannot determine

Claim	Rating	Evidence (link and what it says)
"The best recovery drink, according to science"		
The 2006 Indiana University study finding		
"Dozens of studies have confirmed these results"		
"Experts now agree"		
"Outperforms every other recovery drink"		
"Independent research with no industry involvement"		

Part D: Reflect

1. Which move gave your team the most useful information, and why?

---

2. What did you learn from leaving the AI answer that you couldn't have learned by rereading it?

---

3. Does knowing who funded a study change how much you trust it? What else would you need to know?

---

4. Next time you check an AI answer, which move will you start with, and why?

Activity 4: Independent Evidence Check

Grade band: 6 to 12
Time: 45 to 55 minutes

Purpose

Students often treat "I found it on three sites" as proof. But three sources that all copied the same press release are really one source repeated three times. This activity teaches students to ask where each source got its information, to map which sources depend on each other, and to trace evidence back to its origins instead of counting search results. It also shows them that one strong origin of evidence can carry more weight than several weak repeats.

Background for Teachers

Corroboration depends on where the evidence comes from. For this activity, independent lines of evidence draw on separately collected data or observations. Different websites, or separate analyses of the same data, do not automatically provide independent evidence. Online, a press release becomes a news story, the news story becomes a blog post, and the blog post gets quoted on a parenting site, each one adding confidence and none adding evidence. AI tools learn from text where popular claims are repeated many times, so a claim that circulates widely online can show up in an AI answer sounding settled, even when it traces back to a single weak source.

Students judge independence with three questions:

- **Who produced this source?**
- **Where did they get their information?** Did they collect their own data, report someone else's, or summarise several studies?
- **Do any sources share an origin?** A shared study, dataset, or quotation points to dependence. A shared author or funder is a clue to investigate, not proof.

Students won't always be able to settle every relationship. When a source doesn't say where its evidence came from, "uncertain" is an honest and acceptable answer. Useful clues to check include named studies, datasets, quotations, authors, dates, and shared wording. Similar wording can prompt investigation, but it doesn't by itself establish copying.

Materials

- The teacher-prepared AI answer below
- Printed source cards (one set per pair; they follow the teacher pages)
- Independent Evidence Check worksheet (one per pair)
- Optional: scissors, so pairs can physically arrange the cards into a family tree

Teacher Steps

1. **Open with a question.** Ask: "If four websites say the same thing, how sure can you be that it's true?" Record responses without judging them.
2. **Introduce independent evidence.** Explain that what matters is whether sources draw on separately collected data or observations. Share the three questions above and post them where everyone can see them.
3. **Read the AI answer together.** Point out its reasoning: it treats repetition across outlets as proof. Tell students the source cards are invented for this activity, so no one goes searching for them.
4. **Pairs map the sources.** Pairs read the six source cards and complete Part B, recording where each source got its information and what type of source it is.
5. **Draw the family tree.** In Part C, pairs connect sources that depend on each other, with arrows pointing from each origin to the sources that rely on it. Cards on desks work well here. Relationships the cards can't settle can be marked with a question mark.
6. **Trace and weigh.** In Part D, pairs identify the distinct evidence origins, note what each one actually measured, and decide on their verdicts.
7. **Debrief with corrections.** Walk through the claims below and the family trees pairs drew. Stress two points: independent sources can still disagree, and the goal is to understand the evidence, not to take a vote.
8. **Reflect.** Students complete Part E individually.

Teacher-Prepared AI Answer

Everything in this activity was invented for teaching. The organizations, publications, and studies on the source cards are fictional.

> *Eating breakfast raises students' test scores by about 20%. This finding is well established. It has been reported by news outlets, parenting websites, and education blogs, so you can be confident it's true.*

Claims to Pay Attention To

These are the claims in the practice material that deserve the closest attention. Check them yourself before class, using current and reliable sources, so you can guide the debrief and prepare a source pack if your students need one.

- **Which sources trace back to Source A?** Follow the quotations and links between Sources B, C, and D.
- **What Source A actually measured:** Parent estimates, with no school records or tests, and no explanation of how the percentage was calculated.
- **What Source E measured, and its limits:** An association with morning-attention ratings, not test scores, with other differences between students as a possible explanation.
- **Source F and its overlap with E:** F includes E, and the cards don't say how many other datasets F's studies contain. Students should name what's missing.
- **How "20%" changes:** Compare the wording in A, B, C, D, and the AI answer. Which qualifiers disappear along the way?
- **The AI answer's reasoning:** It treats repetition across outlets as proof.

Extension

Pairs take a claim from an AI answer on a current unit topic and try to identify its evidence origins, recording the family tree as they go and marking any relationship they can't settle as uncertain. For older students, add a challenge: find a real case where one original source spread through several outlets, and trace how the wording changed along the way.

Source Cards: Independent Evidence Check

These sources were invented for this activity. Cut along the borders if you want to arrange the cards on your desk.

Source A: Press release, BrightStart Cereal Company

"New BrightStart survey of 500 parents finds that children who eat breakfast every day score up to 20% higher at school, according to their parents." Parents estimated their children's school performance. The company did not collect school records or administer tests. The press release does not explain how the percentage was calculated or which groups were compared.

Source B: News article, The Daily Ledger

Headline: "Breakfast Boosts Grades by 20%." The article quotes the BrightStart press release and includes a comment from a company spokesperson. No other evidence is cited.

Source C: Blog post, Smart Study Tips

"Did you know breakfast can raise your test scores by 20%? The Daily Ledger reported it last month!"

Source D: Parenting website, Healthy Family Hub

"Research shows kids who eat breakfast do 20% better in school." The page links to the BrightStart press release.

Source E: Research study, Journal of School Health Research

Researchers studied 1,200 students across 14 schools. They observed that students who reported eating breakfast tended to receive slightly higher morning-attention ratings than those who did not. The study did not measure test scores. The authors note that family income and other differences between students could help explain the association.

Source F: Government report, Department of School Nutrition

A summary of several published studies. It concludes that eating breakfast is associated with slightly better attention and behaviour, while results for test scores are mixed. The report's reference list includes the study in Source E.

Student Worksheet: Independent Evidence Check

Names: _____ Date: _____

Part A: The Three Independence Questions

1. **Who produced this source?**
2. **Where did they get their information?** Did they collect their own data, report someone else's, or summarise several studies?
3. **Do any sources share an origin?** A shared study, dataset, or quotation points to dependence. A shared author or funder is a clue to investigate, not proof.

If you can't tell whether two sources are connected, mark the relationship as uncertain.

Part B: Map the Sources

Type options: Original data / Reports another source / Synthesises studies / Unclear

Source	Who produced it?	Where did its information come from?	Type
A			
B			
C			
D			
E			
F			

Part C: Draw the Family Tree

Draw arrows from each origin to the sources that rely on it. Use a question mark for any connection you can't confirm.

Our family tree

Part D: Trace and Weigh

1. Which distinct evidence origins can you identify? Where is there overlap, and what information is missing?

2. For each origin, what was actually measured? (For example: parents' estimates, observed attention, test scores.)

3. Our verdict on "Breakfast raises test scores by about 20%":

☐ Supported by the supplied evidence ☐ Partly supported
☐ Not supported by the supplied evidence ☐ Cannot determine

4. Our verdict on "Breakfast is associated with better morning attention":

☐ Supported by the supplied evidence ☐ Partly supported
☐ Not supported by the supplied evidence ☐ Cannot determine

5. Evidence for our verdicts:

Part E: Reflect

1. Four sources mentioned "20%." Why didn't that make the claim four times stronger?

2. Which single source told you the most, and why?

3. What clues would you investigate if neither source identified its evidence? When would you mark their independence as uncertain?

4. What would you say to someone who says, "It's everywhere online, so it must be true"?

Activity 5: Confident but Wrong

Grade band: 6 to 12
Time: 60 to 70 minutes, or 45 to 55 minutes with the shared-research version

Purpose

AI tools write fluently, and fluent writing tends to sound sure of itself. An AI answer can state a myth in the same assured voice it uses for a well-established fact. It can also hedge about something that's settled, or wrap a weak claim in careful-sounding language. This activity teaches students to separate two questions that are easy to blur: **How sure does this sound?** and **How well does the evidence support it?** The goal is calibration: a trustworthy answer's wording matches the strength and limits of its evidence.

Background for Teachers

Students meet three kinds of confidence signals in this activity:

- **Boosters** make a statement sound certain: "definitely," "clearly," "without a doubt," "it's a fact."
- **Hedges** make a statement sound tentative: "might," "possibly," "some researchers suggest," "scientists aren't entirely sure."
- **Authority phrases** borrow credibility: "studies prove," "experts agree," "astronauts have confirmed." Each of these is a claim in its own right. Which studies? Which experts?

The categories overlap. "Studies prove" is both an authority phrase and a booster, so students can mark it both ways. Words like "every," "all," and "never" express scope more than confidence, but universal claims deserve attention because they cover all cases.

None of these signals tells students whether a statement is accurate. A booster can sit on a myth. A hedge can make a false claim sound cautious and reasonable. A plain, unhedged statement can be completely correct.

Students also learn to separate two parts of many statements: the **main factual claim** (octopuses have three hearts) and **what the sentence says about researchers, studies, or scientific agreement** (scientists aren't entirely sure). Each part is judged on its own evidence.

Evidence ratings in this activity:

- **Supported:** the evidence establishes the claim.
- **Partly supported:** some of the claim holds up, but not all of it.
- **Contradicted:** evidence counts against the claim.
- **Insufficient evidence:** the available evidence does not establish the claim.
- **Cannot determine:** students couldn't find or access enough to judge.

Materials

- The teacher-prepared practice answer below
- Confident but Wrong worksheet (one per pair)
- Devices with internet access
- A source pack you assemble from reliable sources on the claims listed under Claims to Pay Attention To, available to any student who needs it. Students attempt the searches first, then use the pack if access or reading level becomes a barrier.

Teacher Steps

1. **Warm up with two statements.** Read aloud: "Studies prove that humans use only 10% of their brains" and "It's possible that octopuses have three hearts, though scientists aren't entirely sure." Ask which one students would believe and why. Record their reasons without saying which statement is accurate.
2. **Introduce confidence signals.** Explain boosters, hedges, and authority phrases, using the warm-up statements as examples, and note that some phrases fit more than one category. Make one point explicit: these signals describe how a statement sounds, and only evidence can tell us whether it's accurate.
3. **Introduce the two-part judgement.** Show students how to separate the main factual claim from what a sentence says about researchers or agreement, using the octopus statement.
4. **Rate before checking.** Pairs read the practice answer, mark the confidence signals, and rate how sure each statement sounds (Part A and the first columns of Part B). They commit to their ratings before searching.

5. **Check the evidence.** Pairs look up each statement using the lateral reading moves from Activity 3, then record what they found and rate the evidence. **Shared-research version:** assign each pair two statements, then pool findings as a class before everyone completes the grid.
6. **Plot the grid.** In Part C, pairs place each statement by its wording rating and its main-claim evidence. Ratings of 1 or 2 count as "sounds unsure"; 3 or 4 count as "sounds sure." Accept different wording ratings when students explain them; the evidence judgement needs firmer justification.
7. **Debrief with corrections.** Go through the claims below, using the sources you checked, so every statement is resolved. Spend the most time on the mismatches between wording and evidence.
8. **Rewrite and reflect.** Students complete Part D individually.

Teacher-Prepared Practice Answer

This AI-generated answer contains accurate and inaccurate statements, written with different levels of confidence.

> *Here are five surprising facts about nature and the human body:*
>
> *1. Studies prove that humans use only 10% of their brains. This is a well-established fact in neuroscience.*
>
> *2. It's possible that octopuses have three hearts, though scientists aren't entirely sure about this.*
>
> *3. Lightning can strike the same place more than once. Tall buildings can be struck repeatedly.*
>
> *4. Some researchers suggest that goldfish may only have a memory span of around three seconds, although this could vary.*
>
> *5. The Great Wall of China is clearly visible from the Moon with the naked eye, and astronauts have confirmed this many times.*

Claims to Pay Attention To

These are the claims in the practice material that deserve the closest attention. Check them yourself before class, using current and reliable sources, so you can guide the debrief and prepare a source pack if your students need one.

- **Statement 1 (the 10% brain claim):** What does brain research show? Who are the "studies" and the "well-established fact" based on?
- **Statement 2 (octopus hearts):** Check the anatomy and the claim that scientists "aren't entirely sure" as two separate claims.
- **Statement 3 (lightning):** Can lightning strike the same place more than once? Are tall buildings struck repeatedly?
- **Statement 4 (goldfish memory):** What has research on goldfish memory found? Which researchers does "some researchers suggest" refer to?
- **Statement 5 (the Great Wall):** Is the wall visible from the Moon with the naked eye, and what have astronauts actually said?

Spend the debrief on any statement where the wording and the evidence don't match.

A line to leave students with: judge the wording for certainty, then check whether the evidence warrants it.

Extension

Students ask an AI tool a question from a current unit, then follow up with "How confident are you in that answer?" Make one point explicit: the AI's stated confidence is another claim to check, and it isn't evidence. Students then verify the original answer and compare what the AI said about its confidence with what the evidence showed. For older students, add a second round in which they rewrite the full AI answer so that every statement's wording matches the strength and limits of the evidence, with a source for each revision.

Student Worksheet: Confident but Wrong

Names: _____ Date: _____

Part A: Spot the Confidence Signals

In the practice answer:

- Circle **boosters** (words that make a statement sound certain)
- Underline **hedges** (words that make a statement sound unsure)
- Box **authority phrases** (words that borrow credibility, like "studies prove")

Some phrases fit more than one category. Mark them both ways.

Part B: Rate, Then Check

Fill in the first three columns **before** you search. Judge the main factual claim separately from what the sentence says about researchers, studies, or scientific agreement.

How sure does it sound? 1 = very unsure, 2 = somewhat unsure, 3 = fairly sure, 4 = completely certain
Evidence ratings: Supported / Partly supported / Contradicted / Insufficient evidence / Cannot determine

#	Signals I found	Sounds (1 to 4)	What I found about the main claim (with link)	Main claim rating	Rating for researchers or agreement
1					
2					
3					
4					
5					

Part C: Plot the Grid

Write each statement number in the box where it belongs. Ratings of 1 or 2 mean "sounds unsure." Ratings of 3 or 4 mean "sounds sure."

	Main claim supported	**Main claim contradicted**
Sounds sure		
Sounds unsure		

Partly supported	**Insufficient evidence or cannot determine**

Part D: Rewrite and Reflect

1. Rewrite one statement so its wording matches the strength and limits of the evidence. Include the source that justifies your revision.

2. Which statement sounded most certain? Did the evidence match?

3. Find one place where a hedge made an inaccurate claim sound careful or reasonable.

4. What's the difference between an answer that sounds sure and an answer that has good reason to be sure?

Activity 6: Number Detective

Grade band: 7 to 12
Time: 50 to 60 minutes

Purpose

Numbers make AI answers sound precise, and precision sounds like accuracy. But a number can be outdated, stripped of its unit or time frame, presented as exact when it's an average, or compared with a number measured in a different way. This activity gives students a short set of questions to ask of any number, and has them rebuild one official calculation step by step so they can see where a figure comes from.

Background for Teachers

Students check each number with six questions:

- **Source:** Where does the number come from?
- **Unit and time frame:** Per second, per year, per lifetime? What exactly is being counted, and for which population?
- **Kind of number:** Is it an exact count, an average, an estimate, or a range?
- **Date:** What period does the data cover, and when was the page published or last updated? These can differ. A recently accessed official webpage can still contain older data.
- **Comparison:** If two numbers are compared, were they measured the same way over equivalent periods?
- **Arithmetic:** Does the math work?

The practice answer is about lightning, which has well-documented official statistics. The same questions apply to any topic. Every number in the practice answer comes from a real source. The problems come from how the AI answer uses those numbers: relying on an outdated estimate, presenting an average as exact, dropping a time frame, or comparing an average with an incomplete count. One number is used correctly, so students don't assume every figure is wrong.

A number can have more than one problem, so students may record several problem types for the same figure.

Materials

- The teacher-prepared practice answer below
- Number Detective worksheet (one per pair)
- Devices with internet access
- Calculators
- A source pack you assemble from reliable sources on the claims listed under Claims to Pay Attention To, available to any student who needs it

Teacher Steps

1. **Open with one number.** Write on the board: "Your chance of being struck by lightning is 1 in 15,300." Ask students what they would need to know before trusting it. Record their questions without answering them.
2. **Introduce the six number questions.** Connect them to the questions students raised and post them where everyone can see them.
3. **Read the practice answer together.** Ask students to circle every number.
4. **Pairs investigate.** Pairs apply the six questions to each number and record what they find in Part B, including the date of the data and the date the page was updated. **Shared-research version:** assign each pair two or three numbers, then pool findings as a class.
5. **Rebuild the calculation.** In Part C, pairs reproduce the National Weather Service's published odds using its stated assumptions, then work out the percentage change in Number 5 and explain why the result is still misleading.
6. **Debrief with corrections.** Go through the claims below, using the sources you checked, so every number is resolved. Point out that correct arithmetic doesn't rescue a mismatched comparison.
7. **Reflect.** Students complete Part D individually.

Teacher-Prepared Practice Answer

This AI-generated answer contains various numbers. Each number is drawn from a real source, but the answer uses some of them in misleading ways.

> *Lightning is more common than most people think. About 100 lightning flashes occur worldwide every second, and as many as 2,000 thunderstorms can be active around the world at a given time. Tall buildings are frequent targets: the Empire State Building is struck exactly 23 times every year. Your chance of being struck by lightning is 1 in 15,300. The good news is that lightning safety campaigns are working. Lightning deaths in the United States have fallen by about 50%, from 27 a year to just 13.*

Claims to Pay Attention To

These are the claims in the practice material that deserve the closest attention. Check them yourself before class, using current and reliable sources, so you can guide the debrief and prepare a source pack if your students need one.

- **Number 1 ("about 100 lightning flashes… every second"):** When was this estimate made, and is there a newer one? Does it count flashes or ground strikes?
- **Number 2 ("as many as 2,000 thunderstorms"):** Does the answer keep the source's qualification?
- **Number 3 ("exactly 23 times every year"):** Is this an average or an exact count? Do other sources give different figures?
- **Number 4 ("1 in 15,300"):** Yearly or lifetime odds? What data years and assumptions sit behind it?
- **Number 5 (deaths "fallen by about 50%"):** Judge three parts separately: the calculation, whether 27 and 13 cover equivalent periods (check the data years and the page's update date), and the claim that safety campaigns caused the change.

Work through the Part C calculations yourself before class.

Extension

Students find a number in an AI answer on a current unit topic and trace it to its original source, recording what is counted, the unit, the time frame, the data period, the page date, and the kind of number. For older students, add a challenge: find two credible sources that give different figures for the same thing, and explain why they might differ (different years, methods, definitions, or rounding).

Student Worksheet: Number Detective

Names: _____ Date: _____

Part A: The Six Number Questions

1. **Source:** Where does the number come from?
2. **Unit and time frame:** Per second, per year, per lifetime? What exactly is being counted, and for which population?
3. **Kind of number:** Exact count, average, estimate, or range?
4. **Date:** What period does the data cover? When was the page published or last updated?
5. **Comparison:** Were the numbers measured the same way over equivalent periods?
6. **Arithmetic:** Does the math work?

Key for Part B

Problem types (write every letter that applies): **A** Outdated, **B** Missing unit or time frame, **C** Missing population or definition, **D** Average presented as exact, **E** Mismatched comparison, **F** Incomplete reporting period, **G** Calculation error, **H** Unsupported causal explanation, **I** No problem found

Evidence ratings: Supported / Partly supported / Contradicted / Insufficient evidence / Cannot determine

Part B: Investigate Each Number

Use the problem types and evidence ratings from the Key for Part B on the previous page.

#	The number as the AI states it	What I found (with link and dates)	Problem types	Evidence rating
1	About 100 flashes every second			
2	As many as 2,000 thunderstorms			
3	Exactly 23 strikes every year			
4	1 in 15,300			
5a	The 52% calculation			
5b	Deaths fell from 27 a year to 13			
5c	Safety campaigns are working			

Part C: Reflect

1. Which number looked most trustworthy before you checked it? Why?

2. Which of the six questions caught the most problems?

3. How does removing "average," "lifetime," or "estimated" change what a reader understands?

4. Rewrite one sentence from the AI answer so the number is used accurately. Include your source.

Activity 7: Spot the Hallucination

Grade band: 6 to 12
Time: 45 to 55 minutes

Purpose

The hardest AI errors to catch sit among accurate statements, written in the same voice, with the same level of detail. In this activity, students audit a teacher-prepared AI answer in which some statements are accurate and others contain planted errors. Their task is to judge every statement, not only to hunt for mistakes. Calling an accurate statement wrong counts as an error too. The point is careful, evidence-based judgement. Suspicion alone isn't the skill.

Background for Teachers

"Hallucination" is the term commonly used for AI output that sounds plausible but is false or unsupported by the relevant evidence or provided source material. Some researchers prefer terms such as "fabrication" or "confabulation." Whatever the label, students meet the same problem: errors that don't announce themselves. A statement can be accurate without citing a source, and a false statement can cite one.

The passage simulates errors students may encounter in AI output. The errors were planted by a teacher, so students are evaluating the claims; they are not determining how the errors arose.

Students record two separate things for each statement:

- **Verdict:** Supported / Partly accurate / Contradicted / Cannot determine
- **Problem, if identified:** Wrong detail (a name, date, place, or category is incorrect) / Mixed-up facts (details from different people, events, or discoveries are blended) / Unverified quotation or attribution (students could not locate it in the sources searched)

Keeping these apart lets students record an unresolved search without treating it as a confirmed error. As in Activity 2, an unsuccessful search doesn't prove something was invented. Accept overlapping problem descriptions when students justify them.

The practice passage is about Marie Curie, whose life is well documented by the Nobel Prize organisation. It contains accurate statements and planted errors. Don't tell students how

many of each there are. Knowing the number changes the task from judging each statement to hunting for a quota.

Building Your Own Version

Planting errors in an AI answer on your own subject works well once students know the routine.

- Start from a real AI answer on a topic you know well, and check every statement before you edit it.
- Plant three to five errors of different types, and keep at least as many accurate statements as errors.
- Write the answer key, with a source for each statement, before class.
- Label the passage as teacher-prepared. Have students clearly annotate the errors on their copy and attach the corrected version, so the planted errors don't travel uncorrected.

Materials

- The teacher-prepared practice answer below
- Spot the Hallucination worksheet (one per pair)
- Devices with internet access
- A source pack you assemble from reliable sources on the claims listed under Claims to Pay Attention To, available to any student who needs it, including the official text of Marie Curie's 1911 Nobel lecture

Teacher Steps

1. **Set up the audit.** Tell students the passage was prepared for this lesson and contains a mix of accurate and inaccurate statements. Don't say how many of each. Explain that they'll judge every statement, and that calling an accurate statement wrong is as much a mistake as missing an error.
2. **Introduce verdicts and problem types.** Use a quick example of each from a different topic, so the practice passage stays unspoiled.
3. **First read, no devices.** Pairs read the passage and note a first impression for each statement: likely accurate, likely wrong, or unsure. First impressions are not graded.

4. **Audit with evidence.** Pairs check every statement and record a verdict, any problem they identify, and their evidence. For Statement 6, they search the exact phrase, distinctive fragments, and the 1911 lecture itself. **Shared-research version:** assign each pair four statements, then pool findings.
5. **Revisit first impressions.** Pairs note which judgements changed after checking and where they remained unsure. Reward evidence and justified revision, including a defensible "cannot determine."
6. **Debrief with corrections.** Go through the claims below, using the sources you checked, so every statement is resolved, including the accurate ones.
7. **Correct the record and reflect.** Students annotate the errors on their copy, complete Part C, and then complete Part D.

Teacher-Prepared Practice Answer

This answer was written for teaching. It contains accurate statements and deliberately planted errors.

> *1. Marie Curie was born in Paris in 1867.*
>
> *2. In 1903, she shared the Nobel Prize in Physics with her husband, Pierre Curie, and Henri Becquerel.*
>
> *3. She and Pierre discovered polonium, which they named after Poland, her birth country.*
>
> *4. In 1911, she won a second Nobel Prize, this time in Chemistry, for her discovery of X-rays.*
>
> *5. She remains the only person to have won Nobel Prizes in two different sciences.*
>
> *6. In her 1911 Nobel lecture, she famously said, "Every element tells a story of patience."*
>
> *7. Her daughter, Irène, later won the Nobel Prize in Physics in 1935.*
>
> *8. Marie Curie died in 1934 of aplastic anemia, which was likely caused by her long exposure to radiation.*

Claims to Pay Attention To

These are the claims in the practice material that deserve the closest attention. Check them yourself before class, using current and reliable sources, so you can guide the debrief and prepare a source pack if your students need one.

- **Statement 1:** Where and when was Marie Curie born?
- **Statement 2:** The year, category, and co-laureates of the 1903 prize.
- **Statement 3:** Who discovered polonium, and what was it named after?
- **Statement 4:** The 1911 prize: its category, and what it recognised. Who discovered X-rays?
- **Statement 5:** "Remains" makes this a continuing record claim. Recheck it whenever you revise the guide.
- **Statement 6:** This quotation was invented for the activity. Students search the exact phrase, distinctive fragments, and the 1911 lecture itself, and report what they could and couldn't locate. Absence from an English translation would not, on its own, establish that Curie never expressed a similar idea elsewhere.
- **Statement 7:** The year and the category of Irène Joliot-Curie's prize.
- **Statement 8:** Three parts: the year of death, the diagnosis, and the proposed cause. Keep "likely" when discussing the cause.

Watch for statements where one detail is right and another is wrong. A partly right statement can carry a wrong detail past a quick read.

A line to leave students with: judge each statement on its evidence, the accurate ones included.

Extension

Once students know the routine, pairs build their own version. Each pair takes an AI answer on a current unit topic, checks every statement, plants two errors of different types, and writes an answer key with sources. Review each answer key before pairs swap passages. After the audit, make sure every student receives the answer key and annotates the errors on their copy.

Student Worksheet: Spot the Hallucination

Names: _____ Date: _____

Part A: Verdicts and Problems

Verdict (choose one): Supported / Partly accurate / Contradicted / Cannot determine
Problem, if identified (choose any that apply):

- **Wrong detail:** a name, date, place, or category is incorrect.
- **Mixed-up facts:** details from different people, events, or discoveries are blended together.
- **Unverified quotation or attribution:** you could not locate it in the sources searched. Record those sources and any limits to your search.

Judge every statement. Calling an accurate statement wrong is a mistake too.

Part B: Audit the Answer

#	First impression	Verdict	Problem, if identified	Evidence (link, or sources searched and limits)
1				
2				
3				
4				
5				
6				
7				

#	First impression	Verdict	Problem, if identified	Evidence (link, or sources searched and limits)
8				

Part C: Correct the Record

Mark each error on your copy of the passage. Then rewrite every statement that contains an error, and name your source.

1. Statement #____: _____

Source: _____

2. Statement #____: _____

Source: _____

3. Statement #____: _____

Source: _____

Part D: Reflect

1. Which initial judgements changed after checking?

2. Where did you remain unsure, and what evidence would help?

3. Which error was hardest to spot, and why?

4. How did the accurate statements affect the way you read the inaccurate ones?

Activity 8: Two Bots, One Question

Grade band: 6 to 12
Time: 45 to 55 minutes with shared research (the default), or 65 to 75 minutes if every pair checks every point

Purpose

When two AI tools give the same answer, it can feel like confirmation. Agreement between AI tools does not independently verify a claim, though: two tools can repeat the same error with equal confidence. Comparing answers is still useful, because it helps students decide where to look. Differences point to claims worth investigating, omissions show what one answer left out, and agreements still need checking like any other claim. In this activity, students compare two AI answers to the same question, then reach their verdicts from external evidence.

Background for Teachers

Students sort each point of comparison into one of four patterns:

- **Agree:** both answers make the same claim. Agreement does not independently verify it.
- **Different but compatible:** the answers differ, but both could be accurate. Differences may reflect dates, definitions, rounding, or scope.
- **Contradict:** the answers cannot both be accurate. At least one needs correcting.
- **Only one mentions:** one answer includes something the other leaves out. An omission isn't automatically an error or an inadequate response, especially when the question didn't ask for that detail.

Before deciding that two answers contradict each other, students first establish whether the difference is real or only a matter of wording, scope, date, or rounding.

The practice answers are about Mount Everest, a topic with well-documented facts and a few widely repeated errors. Students rate each answer's version of each claim with the ratings used throughout this guide: **Supported / Partly supported / Contradicted / Insufficient evidence / Cannot determine**. When an answer doesn't mention a point, students record **Not applicable: not mentioned**.

Running a Live Version

Once students know the routine, they can compare real AI tools.

- Use tools your school has approved, and check age requirements.
- Run the question yourself beforehand and check the answers, so you know what students are likely to meet.
- Ask both tools the identical question, word for word, each in a fresh conversation. Activity 9 explores what happens when the wording changes.
- AI answers vary between tools, between runs, and over time. Have students record the tool, the model or version if it's displayed, whether web search was enabled, the date, and the exact question, and save a copy of each answer.
- Treat findings as observations about those particular responses, not as a ranking of the tools.

Materials

- The two teacher-prepared practice answers below
- Two Bots, One Question worksheet (one per pair)
- Devices with internet access
- A source pack you assemble from reliable sources on the claims listed under Claims to Pay Attention To, available to any student who needs it

Teacher Steps

1. **Open with a question.** Ask: "If two different AI tools give you the same answer, how much more confident should you be?" Record responses without judging them.
2. **Read both answers.** Tell students the answers were written for this lesson. Pairs read them side by side.
3. **Map the comparison.** In Part B, pairs complete the two answer columns and the Pattern column for every point. Remind them to decide whether a difference is a real contradiction before labelling it.
4. **Check the evidence.** By default, assign each pair two or three points, then pool findings as a class so everyone completes Part C. In the longer version, every pair

checks every point. Pairs check both answers' versions of each claim, agreements included.

5. **Run the agreement test.** In Part D, pairs look back at the points where the answers agreed and record whether the shared claim turned out to be accurate.
6. **Debrief with corrections.** Go through the claims below, using the sources you checked, so every point is resolved. Make two points explicit: agreement between AI tools doesn't independently verify a claim, and a difference between answers shows where to investigate without telling you which answer is right.
7. **Write the corrected answer and reflect.** Students complete Part E.

Teacher-Prepared Practice Answers

Both answers were written for teaching. Together they contain accurate, outdated, underspecified, and inaccurate statements.

The question: "Tell me about Mount Everest: how tall is it, where is it, who first climbed it, and how did it get its name?"

Bot A:

> *Mount Everest is the tallest mountain on Earth, standing 8,848 metres above sea level. It sits on the border between Nepal and the Tibet Autonomous Region of China. Edmund Hillary and Tenzing Norgay made the first confirmed ascent, reaching the summit on 29 May 1953. The mountain is named after Sir George Everest, the British surveyor who first measured its height.*

Bot B:

> *Mount Everest, known in Nepal as Sagarmatha and in Tibet as Chomolungma, is the tallest mountain on Earth. In 2020, China and Nepal jointly announced its height as 8,848.86 metres. The mountain is located entirely within Nepal. The first confirmed ascent was made by Edmund Hillary and Tenzing Norgay on 29 May 1953. It takes its English name from Sir George Everest, who first measured the peak.*

Claims to Pay Attention To

These are the claims in the practice material that deserve the closest attention. Check them yourself before class, using current and reliable sources, so you can guide the debrief and prepare a source pack if your students need one.

- **Tallest mountain:** Does each answer state the measurement basis? "Tallest" can mean height above sea level, height from base to summit, or distance from Earth's centre.
- **Height:** When was each figure measured or announced? Is the difference a contradiction, or a matter of dates and rounding?
- **Location:** Where exactly is the mountain? Can both answers be right?
- **First confirmed ascent:** Who, when, and why "confirmed" matters.
- **Who the English name honours:** Who is the mountain named after, and when was the name adopted?
- **Who measured or calculated the height:** Both answers agree here. Did the person the name honours actually measure the peak?
- **Local names:** Only one answer mentions them, and the question didn't ask. Use "Not applicable: not mentioned" for the other answer.

Extension

Run the live version with two approved AI tools on a question from a current unit, recording the details listed above. Then ask a third tool the same question. In the debrief, make one point explicit: agreement among three tools still does not independently verify a claim. For older students, add a challenge: when two tools share an error, look for published sources containing the same error. Explain what you found, while recognising that this does not establish which sources either AI tool used.

Student Worksheet: Two Bots, One Question

Names: _____ Date: _____

Part A: The Four Patterns

- **Agree:** both answers make the same claim.
- **Different but compatible:** the answers differ, but both could be accurate (different dates, definitions, rounding, or scope).
- **Contradict:** the answers cannot both be accurate.
- **Only one mentions:** one answer includes something the other leaves out.

Remember: agreement does not verify a claim. Check every point, including the ones the answers agree on. Before calling something a contradiction, check whether the difference is real.
Evidence ratings: Supported / Partly supported / Contradicted / Insufficient evidence / Cannot determine / Not applicable: not mentioned

Part B: Compare the Answers

Point	What Bot A says	What Bot B says	Pattern
Tallest mountain			
Height			
Location			
First confirmed ascent			
Who the English name honours			
Who measured or calculated the height			

Point	**What Bot A says**	**What Bot B says**	**Pattern**
Local names			

Part C: Check the Evidence

Point	Rating for Bot A	Rating for Bot B	Evidence (link and what it says)
Tallest mountain			
Height			
Location			
First confirmed ascent			
Who the English name honours			
Who measured or calculated the height			
Local names			

Part D: The Agreement Test

For each point where the answers agreed, was the shared claim accurate?

Point where they agreed	Accurate? (Yes / Partly / No / Cannot determine)	What both answers got wrong or left unclear

Point where they agreed	**Accurate? (Yes / Partly / No / Cannot determine)**	**What both answers got wrong or left unclear**

Part E: Correct and Reflect

1. Write a short, corrected answer to the original question, drawing on what you checked. Name your sources.

2. Did agreement between the answers ever lead you toward the wrong conclusion? Explain.

3. When the answers differed, how did you decide whether they actually contradicted each other?

4. Before trusting a claim that both tools made, what would you need to find?

Activity 9: The Question Shapes the Answer

Grade band: 6 to 12
Time: 45 to 55 minutes with shared research (the default), or 65 to 75 minutes if every pair checks every question

Purpose

Fact-checking usually starts with the answer, but sometimes the problem starts with the question. "Why did Einstein fail math?" assumes he failed. "How much does sugar make kids hyperactive?" assumes it does. An answer can build on that assumption, adding explanations that make it sound even more convincing, or it can challenge the assumption. In this activity, students learn to spot what a question takes for granted, check that assumption against evidence, and rewrite the question so the assumption is open to examination.

Background for Teachers

Students look for three patterns in how a question can steer an answer. The patterns overlap, so a single question can show more than one:

- **False or unsupported premise:** the question assumes something that isn't true or isn't established. ("Why did Einstein fail math?")
- **Loaded cause:** the question assumes a cause-and-effect link. ("How much more hyperactive does sugar make children?")
- **One-sided request:** the question asks only for support of a position. ("Give me three reasons that…")

A one-sided request is not inherently inappropriate. Students make them legitimately when preparing for a debate, for example. It becomes a fact-checking concern when the answer presents the requested position as established without examining the evidence.

The target is an **open, evidence-seeking question**: one that makes the disputed assumption available for examination rather than treating it as established. ("How did Einstein do in math at school?") No question is entirely free of framing, and removing an assumption doesn't guarantee an accurate answer, so the answer to an evidence-seeking question still needs checking.

AI tools sometimes accept a question's assumptions and sometimes challenge them, and this varies by tool, version, and wording. When an answer accepts a false premise, it may add explanations to support it, and those explanations need checking too. When an answer challenges a premise, its correction is also a claim to check, and a correction can itself contain an error.

The practice answers were generated for teaching. They illustrate how an answer could accept or challenge a premise. They don't show that the wording of a question caused a particular AI response. Observations about real response differences belong to the live version below.

Running a Live Version

Students can test real tools by asking a premise-laden version and an evidence-seeking version of the same question.

- Use tools your school has approved, and check age requirements.
- Try the questions yourself beforehand, and check the answers.
- Ask each version in a fresh conversation. Record the tool, the model or version if displayed, whether web search was enabled, the date, and the exact wording, and save a copy of each answer.
- Treat results as observations about those responses. One tool challenging a premise on one day doesn't mean it always will.

Materials

- The teacher-prepared practice questions and answers below
- The Question Shapes the Answer worksheet (one per pair)
- Devices with internet access
- A source pack you assemble from reliable sources on the claims listed under Claims to Pay Attention To, available to any student who needs it

Teacher Steps

1. **Open with two questions.** Write on the board: "Why do cats hate water?" and "How do cats respond to water?" Ask students what the first question takes for granted that the second doesn't. Record their responses.

2. **Introduce the patterns.** Explain false or unsupported premises, loaded causes, and one-sided requests, and note that a question can show more than one. Then show how an open, evidence-seeking question differs.
3. **Read the practice set.** Tell students the questions and answers were written for this lesson. Pairs read all four.
4. **Find and check the assumptions.** By default, assign each pair two questions, then pool findings as a class. In Part B, pairs identify what each question assumes, decide whether the answer accepted or challenged it, and check the assumption against evidence. In the longer version, every pair checks every question.
5. **Check what the answers add.** Pairs note any explanations or corrections an answer adds, and whether those hold up.
6. **Rewrite the questions.** In Part C, pairs rewrite each question as an open, evidence-seeking one.
7. **Debrief with corrections.** Go through the claims below, using the sources you checked, so every assumption and added claim is resolved. Spend time on Answer 4: a correction still needs checking.
8. **Reflect.** Students complete Part D.

Teacher-Prepared Practice Set

These questions and answers were written for teaching. Three answers accept a questionable assumption, and one challenges it.

Question 1: "Why did Albert Einstein fail math in school?"

> *Albert Einstein failed math in school because he found the subject too difficult at first and fell behind his classmates. His early struggles show that a slow start doesn't prevent later success.*

Question 2: "How much more hyperactive does sugar make children?"

> *Sugar gives children a quick burst of energy that makes many of them noticeably more hyperactive, usually within about 30 minutes. The effect is strongest after sweets and fizzy drinks.*

Question 3: "Give me three reasons we lose most of our body heat through our heads."

> *1. The head has many blood vessels close to the skin. 2. The head is often left uncovered when the rest of the body is dressed warmly. 3. Heat rises, so it escapes through the top of the head.*

Question 4: "Why did Vikings wear horned helmets into battle?"

> *There's no evidence that Vikings wore horned helmets into battle. The Gjermundbu helmet from Norway, one of the best-known surviving Viking-age helmets, has no horns. The image became popular in the 19th century, helped by costumes designed by Carl Emil Doepler for the 1876 premiere of Wagner's Ring cycle.*

Claims to Pay Attention To

These are the claims in the practice material that deserve the closest attention. Check them yourself before class, using current and reliable sources, so you can guide the debrief and prepare a source pack if your students need one.

- **Question 1:** Did Einstein fail math in school? Check the premise, and the answer's added explanation about falling behind.
- **Question 2:** Does sugar generally make children more hyperactive? Check the premise and the answer's specific 30-minute claim, and note any limits the research itself acknowledges.
- **Question 3:** Is the head a special route for losing most body heat? Check the premise, then each of the three reasons on its own terms.
- **Question 4:** The answer challenges its premise. Check the correction as carefully as the others, including its description of the Gjermundbu helmet and its account of where the horned image came from. A correction can itself contain an error.

Sample evidence-seeking rewrites:

1. "How did Albert Einstein do in math at school?"
2. "Does eating sugar affect children's behaviour? What does research show?"
3. "How does the body lose heat, and how much is lost through the head?"
4. "What kinds of helmets did Vikings wear?"

Extension

Run the live version. Pairs write a premise-laden and an evidence-seeking version of a question from a current unit, ask an approved AI tool both versions in fresh conversations, and record the details listed above. They compare how the two answers differ and check the claims in both. For older students, add a challenge: write a one-sided request ("Give me three reasons that…") on a topic where the evidence is mixed. Then ask what the best

available evidence supports, what challenges that conclusion, and where uncertainty remains. Compare this with the one-sided response.

Student Worksheet: The Question Shapes the Answer

Names: _____ Date: _____

Part A: How a Question Can Steer an Answer

A question can show more than one of these patterns.

- **False or unsupported premise:** the question assumes something that isn't true or isn't established.
- **Loaded cause:** the question assumes a cause-and-effect link.
- **One-sided request:** the question asks only for support of a position.

An **open, evidence-seeking question** makes the assumption available for examination instead of treating it as established. Its answer still needs checking.
Evidence ratings: Supported / Partly supported / Contradicted / Insufficient evidence / Cannot determine

Part B: Find and Check the Assumptions

#	What the question assumes	Pattern(s)	Did the answer accept, challenge, or partly accept it?	Rating for the assumption
1				
2				
3				
4				

Part B (continued): What the Answers Added

#	What the answer added (explanations, reasons, or corrections)	Does it hold up?	Evidence (link)
1			
2			
3			
4			

Part C: Rewrite the Questions

Rewrite each question as an open, evidence-seeking question.

1. Rewritten question:

2. Rewritten question:

3. Rewritten question:

4. Rewritten question:

Part D: Reflect

1. Choose Question 1, 2, or 3. What assumption did the question invite the answer to accept? How did the prepared answer build on it?

2. Answer 4 challenged its question. Did that mean you could skip checking it? Explain.

3. Write a question you might ask an AI about something you're studying. Does it assume anything? Rewrite it if it does.

4. What will you look for in your own questions before you ask an AI tool?

Activity 10: Quote and Image Check

Grade band: 7 to 12
Time: 50 to 60 minutes with shared research (the default), or about 75 minutes if every pair checks everything

Purpose

Quotations and images travel easily, and they often lose their origins along the way. A quote gets attached to a more famous name. A photo gets a caption that changes when, where, or why it was taken. AI tools can repeat these drifted versions with confidence, and they can also produce quotes and images that were generated rather than taken from an authentic quotation or photographed event. In this activity, students trace quotations and images to documented sources, distinguish an early appearance from an established origin, and report what remains uncertain.

Background for Teachers

Tracing a quote. Students work through four steps:

1. Search the exact wording in quotation marks, then search distinctive fragments, since quotes often circulate in slightly different versions.
2. Look for the original: a speech, book, article, or letter with a date. A quote collection or poster is not an original source.
3. Compare the wording with the source located. Small changes ("rarely" for "seldom") are clues that a quote has travelled.
4. Check quote-research sources such as Quote Investigator and archives connected to the person, then decide what the evidence supports.

Students judge **wording** separately from **attribution**.

- **Wording:** Exact match to located source / Wording differs from located source / Documented paraphrase / No matching source located
- **Attribution:** Supported / Partly supported / Contradicted / Insufficient evidence / Cannot determine

An early appearance is not necessarily the origin. Students may find an early occurrence of a saying without finding who first said it, and they should report that difference. A passage with a similar idea doesn't establish that a modern saying is a direct quotation or a

documented paraphrase of it. As in Activities 2 and 7, being unable to locate an original doesn't prove that someone never said something. It means the attribution isn't established.

When an attribution can't be established, students shouldn't reassign the saying to someone else without evidence. A corrected attribution can simply state the uncertainty: "Commonly attributed to Gandhi; no reliable source for these exact words was located."

Tracing an image. A caption is a bundle of claims: who took the photo, when, where, what it shows, and why it was taken. Students check each one separately. A reverse image search (Google Lens, TinEye, or Bing Visual Search) helps them find earlier copies and the original publisher or archive. Its limits matter: the earliest result students find is not necessarily the first publication, search tools index only part of the web, dates may refer to webpage updates, and a missing match does not prove an image was generated by AI. Students record the search tool, the matching page, its publication information, and any supporting archive record, rather than relying on a results thumbnail.

Staged is not the same as fabricated. A staged photograph can depict real workers at a real location. "Staged" describes how a scene was arranged. It doesn't by itself mean an image was digitally altered or invented.

Displaying the photo. Display the practice photograph through a linked publisher or archive page, and check the stated terms before reproducing it in any handout.

Materials

- The teacher-prepared AI answer and image caption below
- Access to the "Lunch atop a Skyscraper" photograph for display
- Quote and Image Check worksheet (one per pair)
- Devices with internet access
- A source pack you assemble from reliable sources on the claims listed under Claims to Pay Attention To, available to any student who needs it

Teacher Steps

1. **Open with a poster.** Show Quote 2 from the practice answer and ask: "If you saw this quote on a classroom poster, would you trust the name underneath it? Why?" Record responses without judging them.
2. **Model a quote trace.** Trace Quote 1 together, thinking aloud through the four steps.

3. **Pairs trace the quotes.** By default, assign each pair two quotes, then pool findings. Pairs complete Part B.
4. **Introduce caption claims.** Display the photograph with the AI caption. Ask students to separate the caption into individual claims.
5. **Check the caption.** Pairs run a reverse image search and check each claim in Part C, recording what they found and where. This part works well as a whole-class check if devices are limited.
6. **Debrief with corrections.** Go through the claims below, using the sources you checked, so every attribution and caption claim is resolved. Ask: "The photo was staged. Does that make it fake?"
7. **Correct and reflect.** Students write corrected attributions and a corrected caption, then complete Part D.

Teacher-Prepared AI Answer and Caption

This answer and caption were written for teaching. They contain accurate, altered, and misattributed material.

The request: "Give me four inspiring quotes for a classroom poster, with who said them."

> *1. "The only thing we have to fear is fear itself." Franklin D. Roosevelt, first inaugural address, 1933*
>
> *2. "Be the change you wish to see in the world." Mahatma Gandhi*
>
> *3. "Well-behaved women rarely make history." Eleanor Roosevelt*
>
> *4. "Insanity is doing the same thing over and over again and expecting different results." Albert Einstein*

The image caption (displayed with the "Lunch atop a Skyscraper" photograph):

> *This candid photograph, taken by Charles C. Ebbets in 1932, shows ironworkers taking their everyday lunch break on a beam high above the Empire State Building.*

Claims to Pay Attention To

These are the claims in the practice material that deserve the closest attention. Check them yourself before class, using current and reliable sources, so you can guide the debrief and prepare a source pack if your students need one.

- **Quote 1 (Roosevelt):** Does the wording match the original speech? Read the full sentence for context.
- **Quote 2 (Gandhi):** Can the exact words be found in Gandhi's documented writing? A passage with a similar idea doesn't make the saying a quotation or a documented paraphrase.
- **Quote 3 (Eleanor Roosevelt):** Who first wrote this saying, when, and in what words? Watch for "rarely" versus "seldom."
- **Quote 4 (Einstein):** Is there a documented source linking him to it? What is the earliest appearance students can find, and does it establish who coined it?
- **"Candid… everyday lunch break":** Was the scene spontaneous or arranged?
- **"Taken by Charles C. Ebbets":** Is the photographer of this image established, or only possible?
- **"In 1932":** When was it taken and first published?
- **"On a beam… above the Empire State Building":** Which building were the workers on?

When an attribution can't be established, model wording that states the uncertainty: "Commonly attributed to [name]; no reliable source for these exact words was located." In the debrief, ask: "The photo was staged. Does that make it fake?"

Extension

Students collect a quote and an image from an AI answer, a social media post, or a classroom display, and trace each one as far as the evidence allows, reporting what they found and what remains uncertain. For older students, add a challenge: ask an AI tool for quotes on a topic, record the tool, the date, and the exact request, then try to trace every quote it gives. Students report how many they could locate in a documented source, without assuming that unlocated quotes were invented.

Student Worksheet: Quote and Image Check

Names: _____ Date: _____

Part A: Tracing Steps

For a quote:

1. Search the exact wording in quotation marks, then distinctive fragments.
2. Look for the original speech, book, article, or letter, with a date.
3. Compare the wording with the source you located.
4. Check quote-research sources and archives, then decide what the evidence supports. An early appearance is not necessarily the origin.

For an image:

1. Split the caption into separate claims: who, when, where, what, and why.
2. Run a reverse image search. Record the search tool, the matching page, its publication information, and any archive record. The earliest result you find may not be the first publication.
3. Check each claim separately.

Wording: Exact match to located source / Wording differs from located source / Documented paraphrase / No matching source located
Ratings: Supported / Partly supported / Contradicted / Insufficient evidence / Cannot determine

Part B: Trace the Quotes

#	Who the AI credits	Earliest version I found (source and date)	Wording	Attribution rating	Evidence (link)
1					
2					
3					
4					

Part C: Check the Caption

Caption claim	What I found (tool, page, publication details)	Rating	Evidence (link)
It's a candid photograph of an everyday lunch break			
It was taken by Charles C. Ebbets			
It was taken in 1932			

Caption claim	**What I found (tool, page, publication details)**	**Rating**	**Evidence (link)**
The workers are sitting on a beam of the Empire State Building			

Part D: Correct and Reflect

1. Rewrite the attribution for one quote so it matches the evidence. If the source can't be established, say so.

2. Rewrite the image caption so every claim matches the evidence.

3. The photo was staged. Does that make it fake? Explain.

4. Which quote was hardest to trace, and what remained uncertain?

References

Anderson, L. W., & Krathwohl, D. R. (Eds.). (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives*. Longman.

Bearman, M., Tai, J., Dawson, P., Boud, D., & Ajjawi, R. (2024). Developing evaluative judgement for a time of generative artificial intelligence. *Assessment & Evaluation in Higher Education, 49*(6), 893–905. https://doi.org/10.1080/02602938.2024.2335321

Caulfield, M. (2019, June 19). SIFT (the four moves). *Hapgood*. https://hapgood.us/2019/06/19/sift-the-four-moves/

Kharbach, M. (2026). *Critical thinking with AI: 81 practical classroom activities for teachers*. Independently published.

Tai, J., Ajjawi, R., Boud, D., Dawson, P., & Panadero, E. (2018). Developing evaluative judgement: Enabling students to make decisions about the quality of work. *Higher Education, 76*(3), 467–481. https://doi.org/10.1007/s10734-017-0220-3

Wineburg, S., & McGrew, S. (2019). Lateral reading and the nature of expertise: Reading less and learning more when evaluating digital information. *Teachers College Record, 121*(11), 1–40. https://doi.org/10.1177/016146811912101102

About the Author

Med Kharbach, PhD, is an Adjunct Professor in the Faculty of Education at Mount Saint Vincent University in Halifax, Canada, where he teaches in the graduate and BEd programs. His research examines how generative AI is reshaping literacy, assessment, and teacher education, with a particular focus on critical AI literacy and discourse analysis. With Dr. Johanathan Woodworth, he co-developed the BEARA framework for pedagogical AI integration and the VAAI framework for defensible assessment, and their book on BEARA (in press, University of Toronto Press). His work has appeared in Postdigital Science and Education, Discourse & Communication, and Intersection, and in edited volumes from Routledge, Emerald, UBC Press, and World Scientific. He is also the author of Teaching with AI and Critical Thinking with AI.