NO PREP

# AI for Research

## Fact-Checking and Source-Finding

**Includes:** Student reading, case study, differentiated questions, teacher key

Name: Period: Date:

### AI for Research: Fact-Checking and Source-Finding

**The Dangerous Confidence Problem**

AI tools are extremely good at sounding right. When you ask an AI a factual question, it responds in a confident, clear, authoritative tone — the tone of something that knows what it is talking about. This is a significant problem, because AI can be confidently wrong. It can invent statistics. It can cite studies that do not exist. It can get historical dates wrong. It can state consensus positions in fields where there is actually significant debate. And it will do all of this in the same calm, assertive voice it uses when it is completely correct. This phenomenon is sometimes called hallucination — the AI generates content that sounds plausible and true but is fabricated. The term is a bit misleading because it implies something accidental or unusual,

**Name:** **Period:** **Date:**

# AI for Research: Fact-Checking and Source-Finding

### The Dangerous Confidence Problem

AI tools are extremely good at sounding right. When you ask an AI a factual question, it responds in a confident, clear, authoritative tone (the tone of something that knows what it is talking about). This is a significant problem, because AI can be confidently wrong. It can invent statistics. It can cite studies that do not exist. It can get historical dates wrong. It can state consensus positions in fields where there is actually significant debate. And it will do all of this in the same calm, assertive voice it uses when it is completely correct. This phenomenon is sometimes called hallucination: the AI generates content that sounds plausible and true but is fabricated. The term is a bit misleading because it implies something accidental or unusual, when in reality it is a structural feature of how large language models work. They predict what words should come next based on patterns in training data. Sometimes those patterns produce accurate facts. Sometimes they produce plausible-sounding falsehoods. The model cannot distinguish between the two because it has no ground truth to check itself against.

### What AI Research Is Actually Good For

Despite the hallucination problem, AI is genuinely useful in the early stages of research. What it does well is provide a broad orientation to a topic, helping you understand the landscape of a subject before you dig into primary sources. If you are starting a research paper on criminal justice reform and have no idea where to begin, asking AI to outline the five or six major areas of debate is a legitimate and useful starting point. You will get a sense of the terrain: mandatory minimums, recidivism, prison conditions, police reform, bail reform, disproportionate sentencing. You can then go research each of those areas properly using verifiable sources. AI is also useful for explaining concepts and terminology that appear in

your research. If you are reading a scholarly article and encounter a term you do not understand, asking AI to explain it in plain language is faster and often more effective than searching a dictionary. If you are confused about how a specific policy works, AI can give you an accessible explanation. These uses are low-stakes; you are using AI to understand the language of your sources, not to generate the substance of your argument.

**Why You Should Never Cite AI Directly**

There are several reasons why citing AI as a source in academic or professional work is problematic. First, AI outputs are not verifiable: there is no URL you can check, no page number, no publication record. Second, the same question asked to the same AI on different days may produce different answers, which means the source is not stable or consistent. Third, AI can be wrong in ways that are not obvious to a reader, and citing it as a source implies a level of reliability it cannot provide. Most educational institutions have policies against citing AI directly, and for good reason. The purpose of a citation is to let your reader check your source, evaluate its credibility, and see whether it says what you claim it says. AI cannot support any of those functions. If AI tells you a fact, that fact needs to be verified through a citable source before you can use it. Think of AI as a tipster who gives you leads. A good journalist does not print a story based on a tipster's claim alone, they verify it independently before publishing.

**Lateral Reading: The Professional Approach to Verification**

Professional fact-checkers and research librarians use a technique called lateral reading to verify information. Instead of reading deeply into one source to evaluate its credibility, you leave that source immediately and search for what other sources say about it. You are checking the source, not just the content. Applied to AI research, lateral reading means: when AI gives you a claim, you do not spend time evaluating whether the AI's explanation sounds credible. Instead, you immediately open a second tab and search for that claim in databases, news archives, government reports, or academic sources. If you can find multiple independent sources that confirm the claim, it is probably reliable. If you cannot find any source that mentions the statistic AI cited, that is a red flag. If you find sources but they say something different from what AI told you, AI was wrong. Lateral reading is fast

and efficient once you practice it — it takes about 60 seconds to confirm or refute most AI claims.

**Finding Real Sources: What AI Can and Cannot Do**

Some AI tools have the ability to search the web in real time and surface links to actual sources. When AI does this, those links are worth following — they represent real, checkable information. When AI gives you citations or links without web search capability, be skeptical. AI language models can fabricate citations that look real but do not exist. A realistic-sounding author name, journal title, volume number, and page range that produces zero results in a database search is a hallucinated citation. The safest approach is to use AI to generate a list of topic areas and key search terms, then do your actual source-finding in Google Scholar, JSTOR, ProQuest, PubMed, government websites (.gov), or your school or library databases. AI is good at telling you what kinds of sources exist and what to search for. The actual finding and vetting of those sources is still work that requires human judgment and access to real research databases.

**Building a Research Workflow That Works**

A practical research workflow that uses AI responsibly looks like this. Step one: use AI to orient yourself to the topic — ask for major themes, debates, and terminology. Step two: use those themes and terms to build search queries for real databases. Step three: find three to five credible primary sources (peer-reviewed articles, government reports, reputable journalism). Step four: read those sources carefully. Step five: if something you read is unclear, use AI to help explain it. Step six: if AI generates any claims you want to use in your paper, verify them in primary sources before including them. Step seven: cite only sources you have actually accessed and verified. This workflow uses AI where it is genuinely strong (orientation, explanation) and keeps humans in charge of the verification and synthesis that requires judgment and accountability. It is faster than doing all research without AI, and it is more reliable than doing all research through AI alone. One genuinely valuable research application of AI that is often overlooked is using it to help formulate better search queries. If you are researching a topic and not finding good results, try asking AI: 'What are five different academic search terms I could use to find peer-reviewed research about [your topic]?' AI tends to know the specialized vocabulary that academics and researchers use, which is often

different from the everyday language a student would naturally search. Using AI to generate better search vocabulary, then using those terms in Google Scholar or your library database, dramatically improves the quality of sources you find — without ever relying on AI as the source itself.

**Key Information: AI for Research**

| Research Task | AI Role |
|---|---|
| Topic orientation | Strong — AI gives broad overview of themes and debates |
| Terminology explanation | Strong — AI explains jargon from sources in plain language |
| Citing AI as a source | Not appropriate — AI output is unverifiable and unstable |
| Hallucination | AI generates plausible-sounding but fabricated facts or citations |
| Lateral reading | Verifying AI claims by checking independent sources, not just the AI itself |
| Finding credible sources | Use AI for search terms; do actual source-finding in databases |
| Fabricated citations | AI can invent author names, journal titles, and page numbers that do not exist |
| Verification step | Any AI-generated fact used in writing must be confirmed in a primary source |

**Case Study: The Citation That Didn't Exist**

**Tomás is writing a research paper on the effects of social media on teen mental health. He asks AI for statistics and gets several: 'A 2022 study published in the Journal of Adolescent Psychology found that teens who spend more than 3 hours daily on social media are 47% more likely to report symptoms of anxiety.' Tomás includes this statistic in his paper with the citation the AI provided. His teacher flags it. There is no such study. There is no 2022 article with that finding in the Journal of Adolescent Psychology. The journal exists — but the specific study does not. The AI invented a plausible-sounding citation with real-sounding details. Tomás did not know this was possible. He assumed that if AI gave him a specific journal title, volume, and statistic, it had to be real. He goes back and does lateral reading — searching for real studies on social media and teen anxiety in Google Scholar. He finds several actual peer-reviewed studies with real findings he can verify and cite. His paper ends up stronger with real evidence than it would have been with the invented statistic.**

1. What specific error did Tomás make in trusting AI's citation, and what should he have done differently based on the reading's guidance?

2. Ironically, Tomás's paper ended up stronger after the error was caught. What does this suggest about the relationship between using real sources and the quality of research writing?

**Multiple Choice**

1. What does the term 'hallucination' mean when applied to AI, as described in the reading?

a) When AI makes errors only in technical or scientific topics

b) When AI produces an output too long and repetitive to be useful

c) When AI refuses to answer any question because it is unsure of the facts

d) When AI generates plausible-sounding content that is actually fabricated

2. What is lateral reading, as described in the reading?

a) Leaving a source immediately and checking what other independent sources say about the claim

b) Using AI to summarize multiple sources that are positioned side-by-side for direct comparison

c) Reading multiple chapters of the same source to find contradictions within it

d) Reading across multiple drafts of the same document to identify changes

3. According to the reading, what is the safest use of AI in finding research sources?

a) Asking AI to evaluate whether sources you found elsewhere and online are actually credible

b) Asking AI to generate a full bibliography with citations for your paper

c) Using AI-provided citations directly because they include specific journal names

d) Using AI to generate topic areas and search terms, then finding real sources in databases

**Matching**

Write the letter of the correct definition on the line next to each term.

| Term | Definition |
|---|---|
| 1. ___ Hallucination | A) Checking a claim by searching what independent sources say rather than evaluating the original source deeply |
| 2. ___ Lateral reading | B) Confirming an AI-generated claim through a citable, independently accessible source |
| 3. ___ Primary source | C) A reference invented by AI that includes real-sounding details but does not actually exist |
| 4. ___ Fabricated citation | D) When AI generates plausible-sounding but factually false content including invented statistics |
| 5. ___ Verification | E) An original research article, government report, or firsthand account used as direct evidence |

**Fill-in-the-Blank**

1. AI can be confidently wrong — generating plausible-sounding falsehoods — because of a structural feature called _____ .

2. The purpose of a citation is to let your reader check your source and see whether it says what you _____ it says.

3. Step one of a practical research workflow is to use AI to orient yourself to the topic by asking for major themes, debates, and _____, which you then turn into search queries for real databases.

**Short Answer**

Answer each question in 2-4 complete sentences.

1. Explain why AI should not be cited as a source in academic research, using at least two specific reasons from the reading.

2. Describe the responsible research workflow outlined in the reading. What role does AI play and at what point does the human researcher take over?

3. Why is the confident tone of AI responses a potential danger for researchers, and what habit should researchers develop in response?