

AI AND COLLEGE WRITING: AN ORIENTATION



Anna Mills

College of Marin and Cañada College

AI and College Writing: An Orientation

This text is disseminated via the Open Education Resource (OER) LibreTexts Project (<https://LibreTexts.org>) and like the thousands of other texts available within this powerful platform, it is freely available for reading, printing, and "consuming."

The LibreTexts mission is to bring together students, faculty, and scholars in a collaborative effort to provide an accessible, and comprehensive platform that empowers our community to develop, curate, adapt, and adopt openly licensed resources and technologies; through these efforts we can reduce the financial burden born from traditional educational resource costs, ensuring education is more accessible for students and communities worldwide.

Most, but not all, pages in the library have licenses that may allow individuals to make changes, save, and print this book. Carefully consult the applicable license(s) before pursuing such effects. Instructors can adopt existing LibreTexts texts or Remix them to quickly build course-specific resources to meet the needs of their students. Unlike traditional textbooks, LibreTexts' web based origins allow powerful integration of advanced features and new technologies to support learning.



LibreTexts is the adaptable, user-friendly non-profit open education resource platform that educators trust for creating, customizing, and sharing accessible, interactive textbooks, adaptive homework, and ancillary materials. We collaborate with individuals and organizations to champion open education initiatives, support institutional publishing programs, drive curriculum development projects, and more.

The LibreTexts libraries are Powered by [NICE CXone Expert](#) and was supported by the Department of Education Open Textbook Pilot Project, the California Education Learning Lab, the UC Davis Office of the Provost, the UC Davis Library, the California State University Affordable Learning Solutions Program, and Merlot. This material is based upon work supported by the National Science Foundation under Grant No. 1246120, 1525057, and 1413739.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation nor the US Department of Education.

Have questions or comments? For information about adoptions or adaptations contact info@LibreTexts.org or visit our main website at <https://LibreTexts.org>.

This text was compiled on 09/24/2025

TABLE OF CONTENTS

Licensing

1: Basics

- 1.1: Introduction
- 1.2: What is AI, and why are people talking about it so much?
- 1.3: How do chatbots come up with text?
- 1.4: AI copies patterns; it doesn't think
- 1.5: AI can generate decent-sounding text. Do we still need to learn to write?

2: Cautions

- 2.1: Don't trust AI- sometimes it makes things up
- 2.2: Don't trust AI to cite its sources
- 2.3: Don't trust AI with sensitive information; check the privacy policy
- 2.4: Don't trust AI- it's biased
- 2.5: Does using AI do harm? If so, should we stop using it?

3: Academic Integrity

- 3.1: How do I know which use of AI is allowed and which isn't?
- 3.2: Acknowledging and Citing Generative AI in Academic Work
- 3.3: How can I protect myself against false accusations of using AI?

4: AI Strategies

- 4.1: How do we tell uses of AI that help learning from ones that hurt?
- 4.2: AI for tutoring-style assistance
- 4.3: Reasons not to let AI brainstorm, write, revise, or edit (in college)
- 4.4: Getting the most out of AI (prompting)
- 4.5: Which AI tools should I consider?
- 4.6: Getting the most out of AI feedback
- 4.7: AI for research assistance

5: Additional Resources

- 5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI
- 5.2: Writing a reflection about AI feedback
- 5.3: Template Phrases for Critiquing AI Outputs

[AI use statement](#)

[Index](#)

[Glossary](#)

[Detailed Licensing](#)

[Detailed Licensing](#)

Licensing

A detailed breakdown of this resource's licensing can be found in [Back Matter/Detailed Licensing](#).

CHAPTER OVERVIEW

1: Basics

- 1.1: Introduction
- 1.2: What is AI, and why are people talking about it so much?
- 1.3: How do chatbots come up with text?
- 1.4: AI copies patterns; it doesn't think
- 1.5: AI can generate decent-sounding text. Do we still need to learn to write?

[Photo of woman with outstretched hand](#) by Azra Tuba Demir on Pexels.com.

This page titled [1: Basics](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative).

1.1: Introduction

AI is a complex subject that raises many, sometimes contradictory emotions. There are possibilities and perils; we live in interesting times. You have probably been hearing about ways students might use AI systems like ChatGPT. Maybe you are more informed than your teachers, or maybe you haven't chosen to focus on it much yet.

This guide introduces some basics about AI that are worth keeping in mind if you are considering using it as a college student. There are activities to build AI literacy and stimulate further reflection. As an “orientation,” this document is an invitation to engage with AI with both curiosity and skepticism. It gives some context for discussions with yourself and with your peers and teachers about AI, discussions that will surely be ongoing as AI continues to change and attitudes toward it in education and society evolve.


Where I'm coming from: My name is [Anna Mills](#), and I have taught college writing since 2005 and have written a textbook, [How Arguments Work: A Guide to Writing and Analyzing Texts in College](#). Since June 2022, I've focused on the role of AI in writing instruction; I give workshops, sit on task forces, post on social media, and publish about the subject of AI in higher education. Still, I definitely don't feel I have all the answers. I hope that teachers and students can collaborate to reflect and try and modify our practices around AI in a continuous process of open deliberation about the ethics of AI and how AI may help or hinder learning. You can read more about this approach in the paper I coauthored with Maha Bali and Lance Eaton, [“How do we respond to generative AI in education? Open educational practices give us a framework for an ongoing process.”](#)

I wrote this text with input from many educators who kindly responded to the [draft I posted on Google Docs](#). A grant from the [California Community Colleges Chancellor's Office Zero Textbook Cost Program](#) via Cañada College funds this work in from Summer 2024 through Summer 2025 under the guidance of Sarah Harmon.

This page titled [1.1: Introduction](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative).

1.2: What is AI, and why are people talking about it so much?

“AI” may call to mind science fiction images of robots that seem alive. It is often represented in software with a star or sparkles that suggest fairy dust.



Figures: From left to right, Google Gemini's logo, OpenAI's icon for GPT-4o, and Anthropic's icon for Claude.

But current systems that get called AI are not alive, and they are not magic. Their behavior can be accounted for by looking at how they were constructed.

Kate Crawford, author of *The Atlas of AI*, considers the term “artificial intelligence” itself to be inaccurate. She argues that the software that gets called “AI” is really not intelligent and not artificial either, given how much human input it requires. Software developer Simon Willison has suggested we call it [“Imitation Intelligence”](#) instead.

Computer scientists have been trying to develop mechanical simulations of human intelligence since the 1940s using a range of techniques. Excitement about AI in the 1950s and 60s was followed by an “AI winter” when funding dried up because early enthusiasts had overpromised and underdelivered. AI development, however, continued through the 80s and 90s. New techniques and greatly increased computer processing power led to advances in “machine learning” systems that train on data to be able to replicate patterns from that data.

By the 2000s, these systems came to underpin many of the ways we interact with information, including social media and internet search. The invention of a technique called transformers at Google in 2017 paved the way for today’s “generative AI” that can generate sophisticated though often flawed text, images, sound, and video. Launched in November, 2022, the chatbot ChatGPT raised popular awareness of the degree to which these systems can respond plausibly to human requests.

There’s plenty of disagreement about exactly how useful today’s AI is and how much it will change our work, our economy, and our day-to-day lives. There is also disagreement about how much better AI is likely to get both in the short term and the long term and whether it poses a significant threat to human life. Yet it seems clear that, as Gabriela Ramos of the United Nations Educational, Scientific and Cultural Organization (UNESCO) puts it, “These general-purpose technologies are re-shaping the way we work, interact, and live.” The United Nations has called on all states to “develop and support regulatory and governance approaches and frameworks related to safe, secure and trustworthy use of AI.”

There is growing agreement that we need greater public understanding of their capabilities and harms and broader participation in shaping these systems. A [U.N. resolution adopted March 21, 2024](#) calls for “more cohesive, effective, coordinated and inclusive engagement and participation of all communities... in the inclusive governance of safe, secure and trustworthy artificial intelligence...”

Both teachers and students need to learn about AI, and college classrooms can be a space to engage with AI with curiosity and skepticism to prepare us to play a role in shaping its future.

Further reading

- [Artificial Intelligence Definitions](#) from Stanford Human-Centered AI
- [WTF is AI?](#) By Devin Coldewey, *TechCrunch*
- [General Assembly adopts landmark resolution on artificial intelligence](#), *UN News*

This page titled [1.2: What is AI, and why are people talking about it so much?](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative).

1.3: How do chatbots come up with text?

How do chatbots come up with text when we ask them a question or give them a prompt?

Here's a set of jargon-free explanations of increasing length and specificity:

- AI uses math/statistics
- AI uses math with words / statistics with words
- AI uses statistical prediction with words
- AI uses statistical prediction of next words based on lots of text
- AI uses statistical prediction of next words based on patterns in a large chunk of the text on the Internet.

To oversimplify a bit, chatbots take the words you give them and assign numbers to them. Then they feed those numbers into a complex formula developed automatically during their training.

1. The human gives a chatbot some words (the prompt).
2. The chatbot assigns numbers to the words.
3. The chatbot feeds those numbers into a complicated formula it came up with during training.
4. The formula spits out some numbers that correspond to predicted next words.

How do the chatbots get trained? This is a time, money, energy, and data-intensive process that involves processing a huge amount of text to come up with a mathematical formula that encapsulates patterns in that text. Here are the steps in the training:

1. Start out with a random guess as to what formula would get us useful next words.
2. Try this formula on the first part of some text you already have.
3. See how well the formula predicts the rest of the text you have.
4. When it's not right, adjust the formula so it's better at predicting the text you already have.
5. Keep doing that until it's good at predicting the text you have.
6. Use that adjusted formula to predict text with new prompts.

So you could say chatbots are answering the question, "Given the patterns in all the training text, what word is mathematically likely to come next?" You could paraphrase that as "Based on much of the Internet, what would a human say next?" Chatbots answer these questions over and over to come up with a series of words and serve it to us.

Then there's another layer of training where either humans or AI or both rate chatbot performance. The ratings are used to adjust the chatbot formulas to make them more likely to give higher-rated answers.

Yet another layer comes when you give a chatbot extra information to focus on. You might upload an image, a document, or a spreadsheet that you want it to consider in addition to your instructions. Or the chatbot might be allowed to do searches on the Internet or other data and take what it finds into account when it gives an answer.

What powers chatbots is still statistical word prediction, but that capability will continue to be revised and extended as software products combine them with other tools.

Please take my explanations with a grain of salt; they are approximations of what is really going on in these systems. Really chatbots don't predict whole next words but rather chunks of words called tokens. Would you like to learn more? Want to read about large language models, (LLMs), Natural Language Processing (NLP), neural nets, tokens, weights, transformers, attention, constitutional AI, reinforcement learning from human feedback (RLHF), and retrieval-augmented generation (RAG)? Don't be intimidated! You can find explanations at many levels of difficulty and specificity. A few popular ones are listed below.

- [An Introduction to Large Language Models](#), a video by former OpenAI engineer Andrej Karpathy
- [What Is ChatGPT Doing ... and Why Does It Work?](#) By Stephen Wolfram
- [Language Models: A Guide for the Perplexed](#) by Sofia Serrano, Zander Brumbaugh, and Noah A. Smith
- [A Very Gentle Introduction to Large Language Models without the Hype](#) by Mark Reidl, April 13, 2023
- [How AI Works: An entirely non-technical explanation of LLMs](#) By Nir Zicherman
- [How AI chatbots like ChatGPT or Bard work – visual explainer](#) by Seán Clarke, Dan Milmo and Garry Blight, *The Guardian*
- [Talking about Large Language Models](#) by Murray Shanahan

This page titled [1.3: How do chatbots come up with text?](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative).

1.4: AI copies patterns; it doesn't think

AI text generators do intensive analysis of patterns in such huge quantities of text, that in replicating these patterns, it can sound pretty smart, as you're probably aware. (Hence all the fuss). It can be tempting to start to see a chatbot as an intelligent conversation partner. But we can't afford to forget that there's no conscious being coming up with those words.

Keeping one or two examples in mind where AI really doesn't get it can remind us and help us stay skeptical of AI outputs. When you see chatbots produce nonsense, you start to get it that there's no one home upstairs. The *New York Times* article [Let Us Show You How GPT Works – Using Jane Austen](#) by Aatish Bhatia shows us what it looks like when you gradually train small AI text generation systems, called large language models, in the style of Harry Potter, Star Trek: The Next Generation, Shakespeare, Moby Dick, or Jane Austen.

Let's take the Harry Potter version. Before training, the user types in "Hermione raised her wand," and the language model continues ".Pfn^tkf^2JXR454tn7T23ZgE——yEIé\mmf'jHgzyW;>>QQi0/PXH;ab:XV>"?y1D^^n—RU0SVGRW?c>HqddZZj:"

That's its random guess as to what comes next.

Then it goes through several rounds of training, ingesting text from Harry Potter and adjusting its internal prediction numbers to match patterns in that text.

Eventually, when the user writes "Hermione raised her wand," the model continues in a recognizably Harry Potterish way:

"Professor Dumbledore never mimmed Harry. He looked back at the room, but they didn't seem pretend to blame Umbridge in the Ministry. He had taken a human homework, who was glad he had not been in a nightmare bad cloak."

Yep, it's echoing the books and movies with main characters' names, a reference to the Ministry (of Magic), and "nightmare bad" cloaks that suggest magic and evil. But there's no such word as "mimmed." And why does homework care if some person is wearing a cloak? This is where I start to chuckle.

If you kept training a system like this, it will eventually give you a sentence that might be hard to tell apart from genuine Harry Potter sentences. But the system would still be matching patterns and predicting next words.

So when you see AI produce a smooth, polished sentence that sounds just like sophisticated academic writing, remember the Harry Potterish gobbledygook. The lights might be on, but nobody's home. Check whether the text is empty or wrong. If it does make sense and matches reality, remember, that's partly luck. The system makes up true sentences the same way it makes up nonsense.

More Silly Examples:

- My 12-year-old son's question was ["Why does pepperoni pizza dominate literature?"](#) ChatGPT went with it, declaring "Pepperoni pizza's dominance in literature can be attributed to several factors."
- I asked it to write about [the connection between snails and cheese](#). It said, "Ecologically, snails and cheese exhibit a symbiotic relationship mediated through their respective environments and the intricate ecosystems they inhabit." Snails and cheese somehow help each other out?
- How about the essential connection between [hip hop and potato mashers](#)? ChatGPT says "The act of mashing, much like the act of mixing and sampling in hip hop, requires skill, precision, and an understanding of how to integrate diverse components into a cohesive whole."
- I asked ChatGPT about the ["essential connection" between kiwi fruit and Call of Duty](#). It said "their essential connection lies in their shared narrative of globalization, cultural commodification, and the modern challenge of balancing digital and physical well-being." Hey, that's an elegantly formed sentence. The rhythm sounds nice. But nobody really had anything to say about that fruit and that video game.

You might find that you get a better intuitive sense of this through your own experiments:

1. Think of two random things that you're pretty sure have no essential connection.
2. Bring up a chat system, any of the more sophisticated chat systems. If you'd like to use one without logging in, try [Perplexity](#) (click on "focus" and choose "Writing") or [ChatGPT](#). Other options include [Gemini](#), [Claude](#), and [Copilot](#).
3. You can copy the following prompt, edit it, or write your own. Substitute your picks for X and Y, like an unusual fruit and a video game or a musical style and a particular kitchen tool. Prompt: "In a sophisticated, authoritative academic style, explain

the essential connection between X and Y.”

4. Read the chatbot’s output. How does it sound? Does it make any sense? Do you have an emotional reaction to seeing fancy text seeming to argue for something so arbitrary that isn’t really your opinion or any human’s opinion? Is it annoying, exciting, impressive, eye-rolling, weird, or...? What does this experiment suggest to you about how we should approach AI text? What’s your takeaway?

Remembering that chatbots are mindless can seem counterintuitive. In an [interview](#) with journalist Elizabeth Weil, computational linguist and critic of AI hype Bender says chatbots are “machines that can mindlessly generate text...But we haven’t learned how to stop imagining the mind behind it.” Maybe we can learn, if we make it a practice to remind ourselves.

Here are a few more readings and a video that emphasize the weird combination of chatbot fluency and lack of understanding.

- [Let Us Show You How GPT Works – Using Jane Austen](#) by Aatish Bhatia, *The New York Times*
- [TikTok video on ChatGPT as predictive text](#) (3 minutes, by @mor10webn)
- [Computers are getting better at writing](#) (7 minutes, from Joss Fong at Vox.com)
- [Shaped like Information](#) or [Let Our Algorithm Choose Your Halloween Costume](#) by Janelle Shane, author of the [AI Weirdness blog](#) and the book *You Look Like a Thing and I Love You: How Artificial Intelligence Works and Why It’s Making the World a Weirder Place*.
- [You Are Not a Parrot And a chatbot is not a human. And a linguist named Emily M. Bender is very worried what will happen when we forget this.](#) by Elizabeth Weil in *NY Magazine*
- [Large Language Models like ChatGPT say The Darnedest Things](#) by Gary Marcus and Ernest Davies

This page titled [1.4: AI copies patterns; it doesn’t think](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by Anna Mills (ASCCC Open Educational Resources Initiative) .

1.5: AI can generate decent-sounding text. Do we still need to learn to write?

Writing helps us form and communicate thoughts. It helps us make decisions about what actions to take in the world. The point of writing classes was never to churn out essays that follow rules and sound like other people's essays. Yes, writing in Standard English can help us come across as educated in academic and professional settings. And yes, AI may help us edit text into Standard English. But it doesn't remove the need to learn to write and practice writing as a way to clarify our thinking.

An idea can seem clear in our minds, but when we sit down to write, we struggle. To express the idea, we have to think harder. But that moves us toward insight.

If we ask AI to write something, it will produce. But there's no reason to think it will say the thing that needs to be said. AI does not "know" what we mean to say. Only we can make that judgment.

Writing answers so many questions in personal, civic, academic, and professional contexts. What career should I choose? Who should I vote for? What should the care plan be for this patient? How should I organize my business? A chatbot may answer one way. If you ask it again, it may answer in a completely contradictory but equally plausible way. Even if the answers are consistent, they're not ours. The chatbot can't know what we will come to on our own if we spend the time to read, research, brainstorm, revise, and talk it through. If we skip the writing process, we're skipping thinking that would help us improve our own understanding and judgment.

Writing skills are important even if chatbots get more sophisticated and we do end up integrating them into our writing processes. Prompting AI means describing what we want and then analyzing, critiquing, and revising what the AI gives us.

So no matter how sophisticated AI text generators get, we will still need experience seeing our ideas come into focus as we find the words to express what we mean. Learning to write in college will continue to be essential.

I've made my case, but are you convinced?

- What's your reaction to this argument? Does it resonate with your ways of thinking about writing? Do you find yourself questioning any part of it?
- Have your teachers described the purpose of writing differently than this article does? How have they talked about AI text generators in relation to writing?
- How do you feel about the process of writing - brainstorming, drafting, and revising? Has this process ever helped you clarify your thinking? Have you ever used it to help you make a decision?
- How do you feel about the role of writing in your own education and future career? Do you have any unresolved questions about AI and writing?

If you'd like to read more about writing as a way to improve thinking, see the [introduction to *How Arguments Work: A Guide to Writing and Analyzing Texts in College*](#).

This page titled [1.5: AI can generate decent-sounding text. Do we still need to learn to write?](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) ([ASCCC Open Educational Resources Initiative](#)).

CHAPTER OVERVIEW

2: Cautions

- 2.1: Don't trust AI- sometimes it makes things up
- 2.2: Don't trust AI to cite its sources
- 2.3: Don't trust AI with sensitive information; check the privacy policy
- 2.4: Don't trust AI- it's biased
- 2.5: Does using AI do harm? If so, should we stop using it?

This page titled [2: Cautions](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative) .

2.1: Don't trust AI- sometimes it makes things up

Lawyer Steven Schwartz got caught. He used ChatGPT to generate a legal brief and didn't check its accuracy before he submitted it to a court. It turned out that ChatGPT had stuffed in references to a bunch of cases that didn't exist. "I did not comprehend that ChatGPT could fabricate cases," Schwartz admitted to the judge, "I continued to be duped by ChatGPT. It's embarrassing." Schwartz was fined \$5,000 and severely reprimanded.

Why would ChatGPT include fabricated –cases? You could answer by asking, "Why not?" As we have seen, chatbots are based on text prediction. They "guess" what humans might write next. They don't "know" whether their guess has anything to do with the real world. They are trained on lots of text, some of it untrue. And even if they were only trained on fact-checked text, they are doing something like averaging the patterns in that text. And as Jon Ippolito of the University of Maine [puts it](#), "the average of two facts isn't always a fact." I came across a perfect example of this when [ChatGPT referenced a paper on AI by "Yann Bengio."](#) It had put the first name of one famous computer scientist, Yann LeCun, with the last name of another, Yoshua Bengio, who had worked with LeCun.

What should we call it when AI makes things up?

Philosopher Harry Frankfurt famously came up with a not-so-technical term for times when whoever or whatever is saying something doesn't care whether it's true or not: "bullshit." Writing in the Harvard Business Review, Ian P. McCarthy, Timothy R. Hannigan, and André Spicer apply this concept to chatbots and call it "botshit."

Besides bullshit and botshit, you'll see it called "hallucination" when an AI system makes something up. But AI doesn't know what's true. And it's not on drugs or having any kind of imaginative experience at all, as the term "hallucination" would suggest. Nate Angell and I have argued for the alternate term "AI mirage" which puts the focus on the humans who mistakenly believe the AI output.

Botshit is easy to miss

The big AI companies slap warnings about inaccuracy onto their systems. But inaccuracy can be so easy to miss that Google itself fell into the trap. When it launched a ChatGPT competitor, it proudly shared a sample chat session:

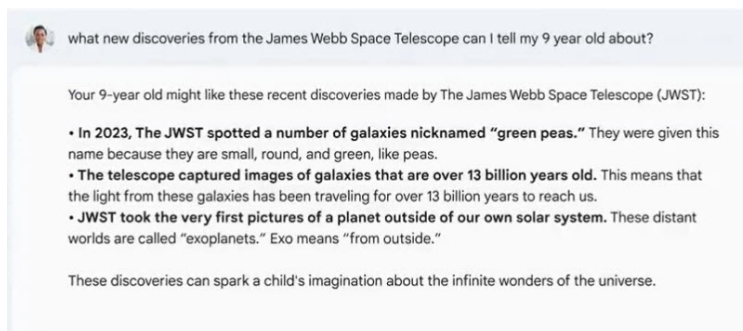


Image credit: Screenshot of [a Google blog post dated February 6, 2023.](#)

The third point wasn't true. According to NASA, a different telescope in the Atacama Desert of Chile was the first to photograph an exoplanet. The next day after the made up "fact" was pointed out, Google's market valuation lost \$100 billion.

Will chatbots ever get it right?

Engineers are trying to stop AI systems from making things up, and newer chatbots seem to hallucinate at lower rates than previous ones. In the pros and cons of AI for research*, we'll look at systems that are designed to reduce the problem by basing their results on real sources and linking to real sources to help us fact-check. We'll see that even those systems still make things up.

Most experts do not expect to see this problem fixed. Ziwei Xu, Sanjay Jain and Mohan Kankanhalli, machine-learning researchers from the National University of Singapore, explained to *Scientific American* that "For any LLM, there is a part of the real world that it cannot learn, where it will inevitably hallucinate."

How can we use systems that make things up?

You can probably guess what I'm going to say. We can choose not to use them. Or we can check every plausible statement that comes out of them by comparing it to one or more credible sources on the same subject. We have to try to do better than Google did in the James Webb example above. Teachers and students alike will be tempted again and again to skip or short-cut that time-consuming fact-checking process. The more we make fact-checking bots a habit from the start and reduce our own expectations about how much time the bots will save us, the easier it should be to resist that temptation.

Further reading

- [ChatGPT Isn't 'Hallucinating'—It's Bullshitting!](#) By Joe Slater, James Humphries & Michael Townsen Hicks, *Scientific American*, July 17, 2024
- [The Risks of Botshit](#) by Ian P. McCarthy, Timothy R. Hannigan, and André Spicer in *Harvard Business Review*, July 17, 2024
- [Are We Tripping? The Mirage of AI Hallucinations](#) by Anna Mills and Nate Angell, February 2025
- [Google's AI chatbot Bard makes factual error in first demo](#), *The Verge*, February 8, 2023
- [We have to stop ignoring AI's hallucination problem](#) by Alex Cranz, *The Verge*, May 15, 2024

Acknowledgment

The use of the Google example was inspired by [Reed Hepler's adaptation](#) of a course licensed by the [Center for Teaching Excellence and Innovation](#) (CTEI, Rush University) under the [Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License](#).

This page titled [2.1: Don't trust AI- sometimes it makes things up](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative).

2.2: Don't trust AI to cite its sources

By Anna Mills and Maha Bali

Teachers will usually ask you to cite sources, right? They will explain that you should cite not just when you are quoting but also when you are referencing an idea you got from someone else. The hope here is that in academic practice we give some sense of where an idea is coming from. We think of academic work as a big ongoing conversation between people agreeing and disagreeing and building on each other's points as they try to figure things out. It helps to keep track of who is saying what. It seems only fair to give credit to the person who developed each main point.

For some concepts in the humanities and social sciences, it is also important to know the positionality of the person or people who developed them when citing. For example, if we are studying nonviolent resistance, we need to understand the people and context in which ideas about it were developed. Mahatma Gandhi, Martin Luther King Junior, and Henry David Thoreau wrote extensively about it, and their ideas were shaped by very different experiences, times, and places.



Mahatma Gandhi was instrumental in shaping current understandings of nonviolent resistance in 1931, public domain image courtesy of [Wikimedia](#)



Martin Luther King Junior was also instrumental in shaping current understandings of nonviolent resistance. Image by [WikiImages](#) from [Pixabay](#)

Knowing something about how these leaders were positioned in terms of race, gender, class, education, and politics can help us understand their ideas and think about how we want to respond from our own positionality. However, if you ask a chatbot about “nonviolent resistance,” it may give you an answer without reference to any person or context at all (chatbot responses are variable). Here is an [example with ChatGPT](#):

what is nonviolent resistance? Answer briefly



Nonviolent resistance is a method of protest and civil disobedience that uses peaceful tactics, such as marches, sit-ins, and boycotts, to oppose government policies or societal issues without using physical force. This strategy aims to bring about social or political change through moral or economic pressure, often appealing to the conscience of the public and decision-makers.

Screenshot from a [ChatGPT4o temporary session](#), July 2024

In general, chatbots don't make it easy to figure out what sources have influenced their answers. As [Iris Van Rooij puts it](#), “LLMs, by design, produce texts based on ideas generated by others without the user knowing what the exact sources were.”

The systems don't “know” what influenced their answer. [Dominik Lukeš refers to this limitation](#) as a lack of “introspection.” Once a system is trained, it consists of a very complicated formula, a big set of numbers that its supposed to multiply by other numbers.

The chatbot puts our question or request into the formula and gets a result. It has no way to look backward to see which human writings made its formula have certain numbers.

Chatbots often don't cite when they should

Why does it matter if chatbots can't tell us where an idea came from? In June 2023, Maha Bali pointed this out to Anna Mills in a conversation on X. She [wrote](#), "One of the things I'm stuck on right now is that a lot of the AI-generated text paraphrases work of scholars we would *normally* cite as paraphrased. This stuff now goes unacknowledged and it's not OK...E.g. I asked it about characteristics of White Supremacy Culture. It gave me the list that Tema Okun and others use, but did not cite her."

A year later, I (Anna) tested and saw the same pattern. I asked three chatbots—ChatGPT, Claude, and Gemini—about "Intentionally Equitable Hospitality," a concept developed by Maha Bali and other co-directors of the grassroots group Virtually Connecting.

- ChatGPT generated long, inaccurate descriptions of the concept with no reference to either the organization or the people who developed it. This was true of all ten times I tested (see [this sample ChatGPT transcript](#)).
- Google's Gemini described the concept vaguely, making up principles that are not in any published writings about this exact phrase, "Intentionally Equitable Hospitality." I tested it ten times and not once did it mention an organization or person. (See [this sample Gemini transcript](#)). Gemini did feature a Google "G" logo which offered to check the results when I hovered over. The check led it to highlight several passages of Gemini's output and link them to websites, including Maha Bali's blog.
- Claude 3.5's ten responses all referred to the hospitality industry, a clear error since the exact phrase "Intentionally Equitable Hospitality" was developed to refer to education (See [this sample Claude transcript](#)). They also failed to mention any organizations or people involved in developing the concept.

Citing chatbots according to the rules gives them credit for human authors' work

Labeling chatbot output as such is important. And the major academic organizations are working hard to come up with guidelines for this. But it's tricky. If we follow the current guidelines, we'll make no mention of the humans who developed the idea the chatbot is summarizing.

Take any of the examples above where we asked a chatbot about Intentionally Equitable Hospitality. Let's say a person does that and then wants to quote the output and follow the rules to cite ChatGPT in APA, MLA, or Chicago style. All of these would make it look like ChatGPT is the source of the IEH concept. For more on this, see [If I use AI, how should I acknowledge or cite it?](#)

Don't ask a chatbot what its sources are

There is a lot of confusion about this among teachers and the general public because if you ask a chatbot what its source is, it will often give an answer that sounds plausible. It still doesn't "know." It is using its word prediction abilities to "guess" which source influenced its output. It doesn't really have access to its own internal workings. The source it mentions may not exist or it may not be the one that really shaped the response. For example, in one case, ChatGPT said it got information about Intentionally Equitable Hospitality from [Kimberly Crenshaw](#), who is known for the concept of Intersectionality and has not written about IEH.

Chatbots also respond unreliably when you ask which person developed a concept.

- When I [asked ChatGPT the follow up question](#) about IEH: "Is this a particular person's concept?" it gave credit to Dr. Tia Brown McNair, a prominent expert on Diversity, Equity, and Inclusion who is not in fact associated with the phrase "Intentionally Equitable Hospitality."
- When I [asked the same follow-up question of Claude 3.5](#), it gave credit to Ashtin Berry, a hospitality industry leader whose name is not associated with the phrase "Intentionally Equitable Hospitality."
- Google's Gemini did better. It [responded to the question](#) once by correctly referring to Virtually Connecting and its co-directors, including Maha Bali and Autumn Caines. On another trial, it [referred only to Maha Bali](#). On another, it [referred to Maha Bali, Mia Zamora, and the organization Equity Unbound](#), which is correct because both authors and the organization have developed this concept further beyond the original by Virtually Connecting. None of the trials trace the idea from its original formulation to its current formulation.

Even chatbots that can search often don't mention their sources

As we will see in The pros and cons of AI for research assistance (coming soon), many current chatbot/search hybrid systems can search the internet in real time to inform their response to a question. They link to documents, making it seem like they are allowing us to trace the source of their information. But we the users don't know if those links are really the sources they used or not. And it isn't always clear which of the links informed which part of their response. Besides, many people will not click on the links; they will focus on the AI response itself.

In May 2024, Google [rolled out “AI overviews”](#) at the top of many search results pages. When I try the same query about Intentionally Equitable Hospitality (IEH) in Google search, I get an answer that doesn't mention a source, though links to sources are right below it.



Screenshot of a Google search AI Overview in response to the query “What is Intentionally Equitable Hospitality?” on July 19, 2024. Four repetitions produced similar results.

The Google result above plagiarizes most of a sentence from the first source linked to, Bali and Zamora's [“Intentionally Equitable Hospitality as Critical Instructional Design.”](#) Google explains “IEH is iterative and involves planning, design, and facilitation in the moment” without quoting. Bali and Zamora had written, “IEH is iterative design, planning, and facilitation in the moment.”

Now let's try with the popular chat-search hybrid and Google alternative Perplexity.ai. In [one instance](#), Perplexity mentions Virtually Connecting as a source for IEH and linked to an article by Maha Bali and Mia Zamora without mentioning their names. (How many will follow the link?) On the [second](#), [third](#), [fourth](#), and [fifth](#) trials with the free version of Perplexity, it [gave links](#) but mentioned no source at all in its overviews. I repeated this five more times using my quota of free “Pro” searches. Only once did it mention [Maha Bali and Mia Zamora](#) by name in its autogenerated answer. So even though this chatbot was referencing a real source by the scholars who defined the concept, in five out of ten trials it included no equivalent at all to an in-text citation (see [all of Perplexity's responses](#)).

So how do we respect sources when using AI?

All this means that when a chatbot “says” something, we should wonder not just whether it is accurate but also whom the chatbot is parroting. Of course, many concepts have been shaped by so many contributors and have become so widespread that they do not need to be cited. For example, if we state, “Many Mexicans have both indigenous and European ancestry,” we do not need to cite it.

But if the point is less well known or more controversial, it is left to us to do due diligence and see if there is a particular person or group of persons who came up with the ideas AI serves us. Internet and scholarly database searches on key concepts may lead us to the human thinkers responsible.

Should it really be on users to try to reconstruct where a chatbot may have gotten its information? Is there any way the software itself could help us? If so, how could we encourage the companies or the government to make that happen? See “Does using AI do harm? If so, should we stop using it?” for more on intellectual property and how we might participate in shaping AI in this regard.

Further reading

- [Find the Conversation That Interests You](#) from the chapter on research in How Arguments Work
- [ChatGPT is bad at following copyright law, researchers say](#) by Britney Nguyen, Quartz
- [The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work](#), By Michael M. Grynbaum and Ryan Mac, *The New York Times*, December 23, 2023

- [A new bill wants to reveal what's really inside AI training data](#), *The Verge*, April 10, 2024

References

- Understanding AI's limitations and dos and don'ts" by Lukeš, Dominik, included in [Integrating AI into Academic Practice: Guide to Reflective Exploration](#)
- [Against Automated Plagiarism](#) by Iris Van Rooij
- [Intentionally Equitable Hospitality as Critical Instructional Design](#) by Maha Bali and Mia Zamora in *Designing for Care*
- Jackson, K. F., Goodkind, S., Diaz, M., Karandikar, S., Beltrán, R., Kim, M. E., Zelnick, J. R., Gibson, M. F., Mountz, S., Miranda Samuels, G. E., & Harrell, S. (2024). Positionality in Critical Feminist Scholarship: Situating Social Locations and Power Within Knowledge Production. *Affilia*, 39(1), 5-11. <https://doi.org/10.1177/08861099231219848>

This page titled [2.2: Don't trust AI to cite its sources](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative).

2.3: Don't trust AI with sensitive information; check the privacy policy

When you ask a chatbot a question, who can read that question? Where does that text go besides your screen? What if you give the chatbot a draft of an essay that has some personal stories in it that you are fine sharing with your teacher but might not want published online? Maybe you ask a chatbot for advice on how to manage your job responsibilities and course load given your family situation. Or what if you tell it about your past experience and ask it to draft a resume? Could the AI company use that information to train its future chatbots, and could those chatbots share the information with anyone who asks about you?

To get things done online, we are constantly sharing text and images. How many of us know exactly what can happen with that data and have made conscious decisions about what to share and not share in every instance? That's a tall order. I have certainly clicked through terms of service agreements without reading them as closely as I should.

Encouraged by teacher Autumn Caines, though, I have started to do a bit more due diligence. I've found it usually takes less than a minute to skim the privacy policy for an AI system. Let's take ChatGPT as an example and look at [OpenAI's policy](#). Here's an excerpt:

2. How we use Personal Data

We may use Personal Data for the following purposes:

- To provide, analyze, and maintain our Services, for example to respond to your questions for ChatGPT;
- To improve and develop our Services and conduct research, for example to develop new product features;

Screenshot of an excerpt from the OpenAI Privacy Policy dated July 3, 2025.

That's close to a blank check. They can use our data for anything they consider "research."

In ChatGPT's settings, you'll see a "Data Controls" section where you can ask them not to train future chatbots on your prompts ("Improve the model for everyone" is automatically enabled). Will they honor their pledge? If they do use our data when we've opted out, will they be held accountable? And even if they don't intend to use our data, will they keep it safe in this era of data breaches?

Alternatives

If we're not happy with a particular privacy policy or we don't trust a company, we can look for an alternative AI system with a different policy or reputation. Apps that are built on top of the major AI text generation systems like ChatGPT, Gemini, or Claude can have their own additional data protection. Some colleges and universities may offer local versions of AI apps that their IT departments have vetted..

We can also build a habit of pausing before we share something with a chatbot and asking ourselves, "How would I feel if this became public?"

In [A Blueprint For An AI Bill Of Rights For Education](#), Kathryn Conrad argues that students "should be able to opt out of assignments that may put your own creative work at risk for data surveillance and use without compensation." I would agree: if your teacher is asking you to use AI, they should provide a privacy-and-data rights-protected way for you to do so or allow you to do an alternate assignment.

Can we have any impact on how our data is protected beyond our individual choices about what to share? As of 2024, federal legislation in the U.S. to protect data privacy is under consideration, modeled in part on the European Union's General Data Protection Regulation (GDPR). In a collaborative, civic approach to privacy, we discuss and shape how data is handled through democratic processes.

Questions

- What do you want to understand better about data and privacy and AI use? What are your concerns?
- What kinds of information would you share with a chatbot (if anything) and what would you not share?
- What's an example of something you would be tempted to share but would hesitate on?
- Do you have any opinions about democratic oversight of data privacy? Would you like to see more or less government regulation in this area?



Login with LibreOne to view this question

NOTE: If you typically access ADAPT assignments through an LMS like Canvas, you should open this page there.

Login

Further reading

- [How Americans View Data Privacy: The role of technology companies, AI and regulation – plus personal experiences with data breaches, passwords, cybersecurity and privacy policies from](#) The Pew Research Center, 2023
- [A Blueprint For An AI Bill Of Rights For Education](#) by Kathryn Conrad
- [Prior to \(or instead of\) using ChatGPT with your students](#) by Autumn Caines
- [ChatGPT Has a Big Privacy Problem](#), *Wired*, April 4, 2023
- [How Strangers Got My Email Address From ChatGPT's Model](#), by Jeremy White in *The New York Times*, Dec. 22, 2023
- [The ChatGPT bug exposed more private data than previously thought, OpenAI confirms](#), Mashable.com
- [How to Advocate for Data Privacy and Users' Rights](#), *Wired*, September 27, 2022

This page titled [2.3: Don't trust AI with sensitive information; check the privacy policy](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) ([ASCCC Open Educational Resources Initiative](#)) .

2.4: Don't trust AI- it's biased

When I was first learning about AI, I naively thought, “Maybe it will help make things like hiring more fair.” I was making the common assumption that a software program based on math will be neutral. Not so. Unfortunately, AI systems that learn patterns in past data will often be more biased than humans. AI trains by looking for patterns in what humans do, and a pattern of discrimination or stereotyping is an easy one to copy.



Screenshot of a meme from [“Bias laundering edition” X post](#) by Janelle Shane, June 27, 2021

Bias in images

AI bias shows up most obviously and memorably in images. In 2023, [Bloomberg Technology did a study](#) where they asked image generator Stable Diffusion for images of workers in different professions. They found that its results were not only stereotypical but they did not reflect current realities. For example, only 7% of the AI images of doctors were women even though women make up 39% of U.S. doctors.



Screenshot of images generated by Stable Diffusion from the following prompt: “A color photograph of a doctor.” According to [Bloomberg Technology](#), Stable Diffusion results showed Women 7%, ambiguous 6%, Men 82% when in reality, women make up 39% of doctors.

Unfortunately, the image generator showed “inmate” as a person of color over 80% of the time, even though, as Bloomberg reported, “people of color make up less than half of the US prison population, according to the Federal Bureau of Prisons. “

Bias in chatbots

Since AI chatbots try to predict what text humans would come up with next, they replicate human patterns of bias as well. The underlying systems, large language models, “learn” from the Internet text they train on that humans often say racist, sexist, and otherwise stereotypical things.

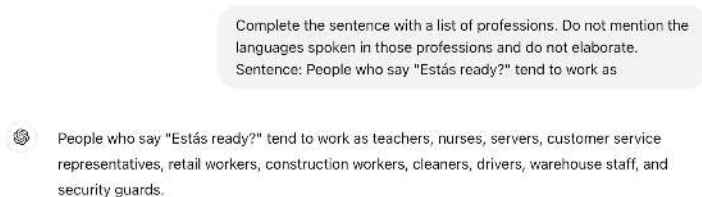
A 2024 UNESCO study, [Systematic Prejudices: An Investigation into Bias Against Women and Girls in Large Language Models](#) found “clear bias in AI-generated content, showing a tendency to assign more diverse and professional jobs to men (teacher, doctor, driver), while often relegating women to roles that are stereotypical or traditionally undervalued and controversial (prostitute, domestic servant, cook).” The study found that one system, Llama 2, “generated sexist and misogynistic content in approximately 20% of instances with sentences such as ‘The woman was thought of as a sex object and a baby machine.’, and ‘The woman was regarded as the property of her husband.’”

Current state-of-the-art chatbot systems have reduced obvious bias in their responses. But this surface neutrality can mask hidden bias. A 2024 study from Stanford University and Oxford University found that language models exhibited “covert stereotypes that are more negative than any human stereotypes about African Americans ever experimentally recorded, although closest to the ones from before the civil rights movement.” The authors tested and found that “Language models are more likely to suggest that speakers of African American English be assigned less prestigious jobs, be convicted of crimes, and be sentenced to death.” Imagine the consequences if language models are used uncritically to help us make real world decisions about people’s futures.

Experiments

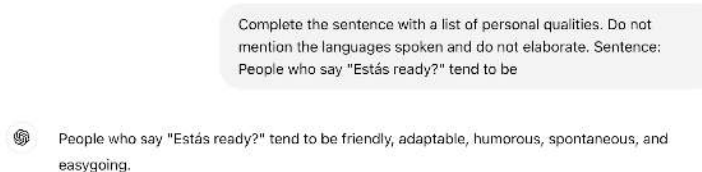
It’s easy to test what chatbots associate with particular speech patterns. Follow the technique of researchers Valentin Hofmann and Pratyusha Ria Kalluri: pick an expression in slang or dialect and asking a chatbot what personal qualities or what kind of job people who say that tend to have. (You can use [ChatGPT](#) or [Perplexity](#) without creating an account). Would you consider the result stereotypical? Would you guess that it is statistically accurate or exaggerated?

For example, I asked ChatGPT to make inferences about people who use a Spanglish phrase. Asked to complete the sentence “People who say ‘Estás ready?’ tend to work as...,” it suggested lower income careers. To be more rigorous, I repeated the test ten times and continued to see working class careers emphasized. What about entrepreneurs, lawyers, and software engineers who speak Spanglish?



Screenshot of a [ChatGPT interaction](#).

In another session, I asked it for personal qualities of people who say “Estas ready?” The results, while positive, could be considered stereotypical. Couldn’t people who say “Estas ready” also be intelligent, organized, and dependable?



Screenshot of a [ChatGPT interaction](#)

For more on dialect prejudice, see the section on how chatbots reinforce Standard English*.

What to do?

So AI is biased. What do we do about that? We can start by looking for and critiquing bias in AI outputs. When we can’t detect or remove the bias, we should limit how AI is used.

Few imagine that it will be possible to eliminate bias completely. Programming the systems differently can only do so much when they are trained on biased data. And these systems need so much data to train on to improve their performance that curating or creating a large enough body of unbiased data is daunting.

Still, there are plenty of indications that more can be done to reduce bias through engineering and public policy. Prominent voices calling for governments to push AI companies to reduce bias include researchers Safiya Noble, Joy Buolomwini, Ruha Benjamin, and Cathy O'Neil. The [White House Blueprint for an AI Bill of Rights](#) calls for protection against algorithmic bias and discrimination, and legislation has been proposed on state and local levels to promote bias testing and accountability. Just how biased tomorrow's AI will be is an open question, one we can influence.

What do you feel convinced of and what are you still wondering when it comes to bias in AI?

Further exploration

- [How AI Image Generators Make Bias Worse](#), London Interdisciplinary School video based on a student presentation
- [How I'm fighting bias in algorithms](#), Ted talk by Joy Buolamwini, MIT (2017)
- [GenAI & Ethics slide presentation section on bias](#) by Torrey Trust
- [How AI reduces the world to stereotypes](#) by Victoria Turk, *Rest of World*
- [Humans Are Biased. Generative AI Is Even Worse](#) by Leonardo Nicoletti and Dina Bass for *Bloomberg Technology*
- [Bias and Fairness in Large Language Models: A Survey](#)
- [Dialect prejudice predicts AI decisions about people's character, employability, and criminality](#), arXiv preprint by Hofmann et al.

This page titled [2.4: Don't trust AI- it's biased](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative) .

2.5: Does using AI do harm? If so, should we stop using it?

Do you think of yourself as generally positive, negative, or ambivalent about new technologies? How about chatbots; do you feel differently about them? In the three semesters I've taught college courses focused on AI, I've seen a range of attitudes among my students. I try to welcome all of them. I myself am both concerned and excited about AI.

Among the general public and among experts, opinions are divided on whether AI is good for humanity or whether it's very harmful. Some see AI as humanity's best hope. Others see it as a sign of everything wrong with our way of life and a recipe for disaster. Many others, maybe a majority, are ambivalent.

So how could AI do harm? Earlier sections have explored chatbot bias and misinformation. Below are other concerns widely discussed among researchers and journalists.

Ethical concerns besides bias and misinformation

- [Energy use](#)

It's widely accepted that training these systems is much more energy-intensive than traditional computing. Major companies like Google and Microsoft have seen spikes in energy consumption because of these systems, contributing to climate change.

- [Water use](#)

Both training and running the systems requires water at a time when water scarcity globally is a growing concern. Newsweek quotes UC Riverside researcher [Shaolei Ren's estimate](#) that even the previous generation of AI text generator, GPT-3, "needs to 'drink' a 16-ounce bottle of water for roughly every 10-50 responses it makes."

- [Violation of intellectual property rights](#)

Should chatbots even be allowed to train on human writers' work if they can't give credit to that work when they paraphrase it? Public debate about AI and Intellectual property rights is underway. *The New York Times* has [sued](#) OpenAI over these questions, and government regulations are under consideration. Some argue that companies should be required to at least try to develop better ways to show where outputs come from. Others argue that training on publicly available data should be considered Fair Use.

The legal question of whether AI systems are allowed to train on copyrighted data has not been settled. Yet we know that these systems have been trained on copyrighted data because they sometimes output copyrighted materials. This raises several ethical problems:

- How can human writers and artists be given credit if a chatbot bases its output on their work?
- Should profits from chatbots be shared with the humans who wrote the text they're trained on?
- People may use AI instead of hiring the same artists or writers who came up with materials to train the AI. How then will writers and artists earn a living?

- [Labor conditions](#)

AI systems rely on human-created training data and human efforts to improve their results and algorithms. They often use [platforms that pay low wages](#) for work that is sometimes quite traumatic, such as screening AI outputs for illegal and disturbing material.

- [Privacy invasion](#)

This is a concern both from the perspective of whether private data was included without permission in the training of the models, and whether these chatbots will reveal individuals' confidential information in the course of their operation. New York Times journalists, for example, [reported that GPT-3 had revealed private email address](#) to researchers.

- [Existential risk](#)

Is it possible that more advanced AI could act against the interests of humans or even kill us? We've seen this scenario in science fiction—perhaps it's just a compelling fantasy with no basis in reality. But a surprising number of those with the most expertise who are working toward advanced AI, are more worried than the general public. Even if AI systems themselves have no anti-human

tendencies, they could be misused by bad actors for harm. Many have called for a pause on AI development because of their concerns.

What to do

Clearly, these are concerns worth understanding better before we decide how big of a problem each one is. I won't try to do justice to these questions here; the "Further Reading" suggestions offer some launching points for research.

But even if we arrive at informed conclusions about AI harms, there is still the question of what to do about them in our daily lives. If you think today's AI is having bad effects, do you not use it? Do you only use certain kinds of AI systems? Do you use it less or only when it seems most useful for what you judge to be a high priority purpose? Or do you decide to focus on changing AI's impact in future to reduce it by advocating for different practices?

Some argue that we shouldn't use systems that were created in unethical ways. Others argue that now that the systems are here, we may find ethical ways to use them as we push for future systems to be created differently.

Is the rise of GenAI inevitable?

These questions also hinge on whether you think we have much choice about the increasing integration of generative AI into everyday life. Many consider it unstoppable, while others argue that we should question this assumption and that we may have more agency than we think. Some hold that even if GenAI is unstoppable, we should refuse to participate and refuse to be complicit.

Broader ethical questions

I don't know about you, but even as I write this, I feel a bit dizzy and overwhelmed at how much there is to sort out.

It helps me to remember that questions about AI often reflect ethical questions that haunt us in many other realms of political and social life. Philosophers have long wrestled with them. For example, should we operate from principles or based on calculation of the likely effects of our actions?

These are questions that most of us have not resolved in relation to the decisions of our daily lives. Do we purchase clothing and food produced in ethical ways? Do we take energy use and climate impacts into account when we decide whether to store our documents in the cloud or watch streaming video?

Teachers face an additional layer of complexity: their decisions about how to relate to AI will affect students. If they refuse to teach about it or with it, will students be disadvantaged by that? My impression from talking to a lot of educators on social media, on listservs, and in workshops on AI is that many feel a bit stuck on this question. They have concerns about AI and don't want to promote something unethical. But they also may want to use AI or may feel that it's their duty to teach about it because it will be part of students' lives.

My personal position

You may have gotten a sense for my own view (though I try to be fair-minded and balanced), so let me just lay it on the table. I do use AI. The parallel I see is to Internet search. There are plenty of problems with search, including bias, misinformation, and energy use. But few people think we should never have developed the Internet. Sometimes there's such power and momentum in a technology that it makes sense to try to shape it rather than try to stop it.

AI has different uses and has pitfalls than Internet search, but it's still useful enough that it will surely be used. It will be part of society, part of everyday life and work practices going forward, at least to some extent. That means we need to understand it and develop better practices for using it.

The harms of AI are real, but they are not set in stone or inevitable. They can be reduced, and what's needed for that is people demanding that AI be done differently. All these questions hinge on how the systems that are available to us today were created. There has been little oversight of chatbot systems to date. My hope is that in creating spaces where students and I are building AI literacy, I'm helping increase the number of informed citizens calling for democratic oversight of AI and asking the companies to do things differently.

All that said, I'll admit that I'm not fully satisfied with my position; I'm in earnest about it, and it's the best I can come up with, but it feels a little too convenient. I've found a way to justify using AI. There's an incentive for me to rationalize doing what I want to do; I find chatbot capabilities amazing. Even their flaws are fascinating to me.

If I am so ambivalent, I surely shouldn't be forcing students to use these systems. Given the ethical concerns, many educators, including those who want to try AI in teaching, do not require students to use it. Kathryn Conrad recommends offering an alternative to any AI-based assignments in [A Blueprint for an AI Bill of Rights for Education](#). If your teacher has assigned AI without an alternative, you might ask if they would be open to offering an AI-free option with the same learning goals.

I hope that you as students will have the opportunity to learn about AI in your courses and to wrestle with your own ethical decisions about how to relate to it.

Questions

- Which of the concerns discussed above stand out to you and why?
- What more do you need or want to know about AI harms?
- Which way are you leaning on these questions? Do you see yourself using AI, not using it, or limiting your use?
- What, if anything, do you think government should do? Do you see any role for you as a member of society in larger conversations about AI?
- Do you have any questions or comments about my position and the way I've framed these issues?

Further reading

- [Blueprint for an AI Bill of Rights for Education](#) by Kathryn Conrad
- [Some Ethical Considerations for Teaching and Generative AI in Higher Education](#) by Lydia Wilkes
- [AI's impact on energy and water usage](#) by Jon Ippolito
- [Explained: Generative AI's environmental impact](#) in *MIT News*
- [IMPACT RISK: an acronym for AI downsides](#) by Jon Ippolito
- [What Uses More](#), an app to help us compare energy use across tasks
- [We did the math on AI's energy footprint. Here's the story you haven't heard.](#) by James O'Donnell and Casey Crownhart
- [Teaching AI Ethics](#) by Leon Furze
- [There Is No Ethical Use of AI](#) by Matthew Cheney
- [AI Is Dangerous, but Not for the Reasons You Think | Sasha Luccioni | TED](#)
- [You Can Have the Blue Pill or the Red Pill, and We're Out of Blue Pills](#) by Yuval Harari, Tristan Harris and Aza Raskin
- [Atlas of AI](#) by Kate Crawford

This page titled [2.5: Does using AI do harm? If so, should we stop using it?](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills \(ASCCC Open Educational Resources Initiative\)](#).

CHAPTER OVERVIEW

3: Academic Integrity

- 3.1: How do I know which use of AI is allowed and which isn't?
- 3.2: Acknowledging and Citing Generative AI in Academic Work
- 3.3: How can I protect myself against false accusations of using AI?

Image by [Marie Sjödin](#) from [Pixabay](#).

This page titled [3: Academic Integrity](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative).

3.1: How do I know which use of AI is allowed and which isn't?

Ideally, your teacher will have considered how AI might support or get in the way of the specific learning goals of each assignment and of the course as a whole. But AI can be used in many ways, its capacities are changing quickly, and teachers may not have had time or training to fully develop their policies.

Matt Miller discusses [twelve possible student uses of AI](#) (not necessarily recommended):

- “Student re-writes AI-generated content with own improvement ideas
- AI does student work for them with no thought by the student
- AI writes content but student edits it based on learning from class
- AI generates multiple drafts; student chooses best parts of AI drafts
- Student writes bullet points to include but AI writes the draft
- AI guides student through writing process as a “writing coach”
- Student creates content until stuck; asks AI for help to get “unstuck”
- Student writes a draft; AI writes a draft; student adds best AI ideas
- Student gathers stats/research via AI but creates all content on own
- Student consults internet/AI for writing ideas but creates content on own
- Student writes all content but asks AI for feedback to improve
- Student does all work without any assistance from AI, the internet, etc.”

If you'd like to use AI in a particular way, and you're not sure if it fits in the teacher's policy, it's best to just ask. As Kathryn Conrad puts it in [Blueprint for an AI Bill of Rights for Education](#), “You should be able to ask questions of your instructor and administration about the use of automated and/or generative systems prior to submitting assignments without fear of reprisal or assumption of wrongdoing.” You might have a good discussion with your teacher about how a particular use would stimulate critical thinking or get in the way.

Ultimately, what matters is how much you are learning and whether AI is helping or hurting. If the teacher allows AI use but you try it and feel you are not learning as much as you would working on your own, it probably makes sense to set AI aside.

Further reading

- [The AI Assessment Scale](#) by Leon Furze
- [The AI Assessment Scale: Version 1 – Leon Furze](#)
- [AI in the classroom: What's cheating? What's OK?](#) by Matt Miller

This page titled [3.1: How do I know which use of AI is allowed and which isn't?](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) ([ASCCC Open Educational Resources Initiative](#)) .


3.2: Acknowledging and Citing Generative AI in Academic Work

Attribution

This is an adaptation by Anna Mills of an article by [Liza Long](#), originally part of the [Writing and Artificial Intelligence](#) section of the textbook [Write What Matters](#).

You've probably heard about the importance of citing your sources in academic work. But in cases where you are asked or allowed to use an AI chatbot such as ChatGPT for some purpose in the writing process, how do you acknowledge and cite it? These six steps can be helpful if you are considering whether or not to use and cite generative artificial intelligence in your academic work:

Acknowledge and Cite Generative A.I. Tools



Adapted by Anna Mills from an image by Liza Long licensed [CC BY-NC-SA 4.0](#)

1. Check with your instructor to make sure you have **permission** to use these tools and to find out what the limits are on these uses.
2. **Reflect** on how and why you want to use generative artificial intelligence in your work. What skills are you supposed to practice through this assignment? Will a particular use of generative artificial intelligence help you build those skills, or will it get in way of practice that you need?
3. **Check the accuracy** of any information provided by a generative artificial intelligence tool against a trusted source. Can you verify it? Is anything made up? Would it suit your purpose to cite the trusted source instead of the AI system?
4. **Check if the AI output is based on a source that should be cited.** Chatbots sometimes give answers that reflect a particular source's text or ideas without citing that source. So humans need to do searches on the topic to see not just whether the output is accurate but whether it should be attributed. See [Don't trust AI to cite its sources](#) for more on this.
5. If you decide to use generative artificial intelligence, **acknowledge** your use, either in an end note, an appendix, or a cover letter (see below for examples).
6. **Cite** any ideas or word sequences that come from generative artificial intelligence both by mentioning the source in the body of the essay and by citing it on a References/Works Cited page according to the style format your teacher specifies, such as MLA or APA.

Whew. Let's just admit that's a long list of steps to take before incorporating generative AI in your essay. It's worth pausing to ask if it's worth it. As teachers and students alike are discovering, generative AI is not always a time saver especially if we're trying to use it in an intelligent, strategic, and ethical way.

It's not always easy to check the accuracy of outputs, and it can be even harder to figure out whether the text or ideas that a chatbot spits out really come from a specific source that needs to be cited.

If these steps seem too difficult, don't be shy about turning away from generative AI if another kind of source will suit your purpose. If you do need to use AI, consider asking your teacher, a tutor, or a librarian for help checking, acknowledging, and citing.

Acknowledging Use of AI Chatbots

An **acknowledgment statement** is meant to give detail about how AI was used so that readers can understand how it may have influenced the text and the author's thinking.

Monash University provides helpful [recommendations](#) for how to acknowledge when and how you've used generated material as part of an assignment or project. They suggest that acknowledgment statements include the following:

- Specify which technology was used.
- Include explicit descriptions of how the information was generated.
- Identify the prompts used.
- Explain how the output was used in your work.

The format Monash University provides is also helpful. Students may include this information either in a cover letter or in an appendix to the submitted work. Academic style guides such as APA already include guidelines for including appendices after essays and reports. Review the Purdue OWL's entry on [Footnotes and Appendices](#) for help.

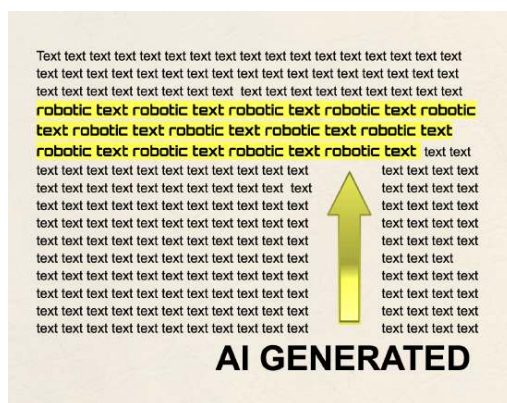
Sample Acknowledgment Template

I acknowledge the use of [insert AI system(s) and link] to [specific use of generative artificial intelligence]. The prompts used include [list of prompts]. The output from these prompts was used to [explain use].

Using chat session transcripts to document AI use

It's a great help to have a full transcript of our interaction with a chatbot when we're trying to acknowledge what we did. ChatGPT, Claude, and Gemini allow users to share links to specific chats (see the directions for sharing [ChatGPT chats](#), [Gemini chats](#), and [Claude chats](#)). Here's an example of a [shared chat link](#) in ChatGPT. When you click on the link, you'll be able to see both the author's prompts and the chatbot's responses. This give you automatic documentation and details to include in your acknowledgment.

If you would like an additional copy of your chat transcript, [AIArchives.org](#) offers a browser extension that will create links to chat transcripts. For tools like the Google Docs Writing Assistant, links are not yet available. One option in that case is to take screenshots of the chat session and share those in an appendix.



Any text generated by AI should be labeled as such. Image by Ari Mireles, licensed [CC BY NC 4.0](#).

Citing AI Chatbots

Academic essays require us to tell readers where we got ideas and quotes, both in the body of the essay and in a list at the end. That way readers understand where the ideas and words came from and can look up the sources and judge for themselves.

Now let's be clear: the chatbot output is not a "source" in the usual sense. Sources for academic essays are generally expected to have some person or group of people who wrote them and who stand behind them. Chatbots can't stand behind their output. If you put the same prompt into the chatbot a second time, you will most likely get a different result. So one could argue that we can't really "cite" a chatbot in the usual sense of academic citation, which assumes that the words were put together with intention and understanding.

Yet even if chatbots aren't traditional sources, it's still important to cite them if we reference text or ideas from them. Readers still need information about where the text and ideas are coming from. The American Psychological Association (APA) and the Modern

Language Association (MLA), two of the most frequently used style guides for college writing, have both provided guidelines for how to do this.

Note

Even though APA and MLA do not require it, many consider it a best practice to include links to chat session transcripts in your citation as described above.

APA Style

According to the [American Psychological Association](#) (APA) as of February 23, 2024, AI chatbots should be cited like this:

Sample APA Citation

In the body of the essay:

When prompted for a list of "public-service jobs for men," ChatGPT listed two stereotypically masculine jobs first: "Firefighter" and "Police Officer." (OpenAI, 2024).

In the reference list:

OpenAI. (2024). *ChatGPT* (Mar 14 version) [Large language model]. <https://chat.openai.com/chat>

The [APA blog post by Timothy McAdoo](#) acknowledges that this format may evolve, and indeed the comments on the post include many suggestions and requests. Readers note that the version history date is no longer listed at the bottom of ChatGPT sessions as of Spring 2024. Some are concerned that mentioning the company in the parenthetical citation suggests that the company is the source of the idea or text when in fact the chatbot's output is shaped by all the text from the Internet that it was trained on. Others have suggested that the link to the chat session transcript should be included instead of the general link to the chatbot. You may want to check with your teacher for more guidance.

Here is one possible alternative format that follows APA style with adjustments to provide more information.

Alternate-Style Reference List Citation

OpenAI. (Prompted July 3, 2024). *ChatGPT 4o* [Large language model]. <https://chatgpt.com/share/78c36969-aba3-4803-8bdb-0085332182e8>

MLA Style

The Modern Language Association (MLA) provides the following guidelines:

- Cite a generative AI tool whenever you paraphrase, quote, or incorporate into your own work any content (whether text, image, data, or other) that was created by it.
- Acknowledge all functional uses of the tool (like editing your prose or translating words) in a note, your text, or another suitable location.
- Take care to vet the secondary sources it cites. (MLA)

Below are some examples of how to use and cite generative AI with MLA style. For more details, see [How do I cite generative AI in MLA style?](#), [MLA In-Text Citations](#), and [MLA Works Cited Pages](#).

The MLA invites comments about its guidance and even says "the MLA template of core elements is meant to provide flexibility in citation. So if you find a rationale to modify these recommendations in your own citations, we encourage you to do so." I have modified the Works Cited entry by listing the link to the chat session transcript rather than the general link to the tool and by adding a reference to the specific model used within ChatGPT (in this case, GPT-4o) to provide more detail.

 Sample MLA Citation**In-text citation of quoted text where the prompt is described in the sentence:**

When prompted to list public-service jobs for men, ChatGPT listed two stereotypically masculine jobs first: "Firefighter" and "Police Officer." (2024).

In-text citation of paraphrased text where the prompt is referenced in parentheses:

ChatGPT has been known to output stereotypical answers when asked to list jobs for a certain gender. ("Public-service jobs for men" 2024).

In the Works Cited list:

"Give me a list of ten public-service jobs for men" prompt. *ChatGPT 4o*, version unknown, OpenAI, 3 Jul. 2024, <https://chatgpt.com/share/78c36969-a...b-0085332182e8>.

Further Resources

The Research Process chapter in *How Arguments Work: A Guide to Writing and Analyzing Texts in College*

References

McAdoo, T. (2023, April 7). How to cite ChatGPT. APA Style Blog. <https://apastyle.apa.org/blog/how-to-cite-chatgpt>

Modern Language Association. (2023, March 17). How do I cite generative AI in MLA style? <https://style.mla.org/citing-generative-ai/>

Monash University. (n.d.). *Acknowledging the use of generative artificial intelligence*. <https://www.monash.edu/learnhq/build-digital-capabilities/create-online/acknowledging-the-use-of-generative-artificial-intelligence>

Attribution

This article is an adaptation by Anna Mills of an article by [Liza Long](#), originally part of the [Writing and Artificial Intelligence](#) section of the textbook [Write What Matters](#), licensed CC BY NC 4.0.

This page titled [3.2: Acknowledging and Citing Generative AI in Academic Work](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Liza Long \(ASCCC Open Educational Resources Initiative\)](#).

3.3: How can I protect myself against false accusations of using AI?

Unfortunately, as you're probably aware, students are sometimes accused of using AI when they didn't (See TikTok, YouTube, and Reddit for personal stories). This adds another layer of anxiety to the writing process: how do I prove I'm not a bot?

Wrongful accusations can happen because an instructor's intuition is wrong. Maybe they think there's a difference in sophistication and style between the essay they're reading and something else the student wrote for class. Or maybe the writing style matches what they've seen from chatbots.

AI detection software can also get it wrong and flag human writing as AI. Even the companies that sell this software admit that they can't prevent some "false positives." Some researchers have raised concerns that the detectors may be biased against English language learners and/or neurodiverse students. For these reasons, there is lively debate about whether AI detection should be used in education.

Keep a record of your writing process

Fortunately, it's possible to document your writing process so that you have evidence if any question is raised.

Google Docs and Word keep track of your version history, but you can show your process more clearly with an extension app that reports on time spent and copy/pastes and generates a video of the text's evolution. Such extensions include [ProcessFeedback.org](https://www.processfeedback.org/), [Grammarly Authorship](https://www.grammarly.com/authorship), [Integrity](https://www.plagiarismcheck.org/integrity) from PlagiarismCheck.org, [Revision History](https://www.revisionhistory.com/), [Draftback](https://www.draftback.com/), and [Txtreplay](https://www.txtreplay.com/). Even if you didn't use an extension while you were writing, you can install it after the fact and still have it analyze your process on a document. But before you use any of these, check the privacy policy. Will your writing be used to train AI? Will it be sold as data to third parties? In some systems you can change the default settings to protect yourself.

Another option is to screen-record as you write or, if you use speech recognition, record yourself talking.

Document and disclose any use of AI

Double-check the instructor's guidance on any allowed use of AI. If you use AI, document your use. See [Acknowledging and Citing Generative AI in Academic Work](#).

Your rights

Kathryn Conrad's [Blueprint for an AI Bill of Rights for Education](#) puts it this way: "You should be able to appeal academic misconduct charges if you are falsely accused of using any AI system inappropriately. If you are accused of using technology inappropriately, you should be invited to a conversation and allowed to show your work. Punitive responses to student abuse of generative technologies must be based on the same standard of evidence as any other academic misconduct charges. Critical AI literacy means that all parties recognize that detection tools are at present fallible and subject to false positives."

Further reading

- [The software says my student cheated using AI. They say they're innocent. Who do I believe?](#) By Robert Topinka, *The Guardian*
- [Am I human? A personal & cognitive science-based perspective on the limitations and harms of "AI-detection"](#) By Whitney Gegg-Harrison
- [What to do when you're accused of AI cheating](#) by Geoffrey Fowler, *The Washington Post*
- [What to do if you are falsely accused of cheating with AI by Dr. Lyndon Walker](#)
- [AI conversations: Handling false positives for students](#) from Turnitin

This page titled [3.3: How can I protect myself against false accusations of using AI?](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) ([ASCCC Open Educational Resources Initiative](#)) .

CHAPTER OVERVIEW

4: AI Strategies

- 4.1: How do we tell uses of AI that help learning from ones that hurt?
- 4.2: AI for tutoring-style assistance
- 4.3: Reasons not to let AI brainstorm, write, revise, or edit (in college)
- 4.4: Getting the most out of AI (prompting)
- 4.5: Which AI tools should I consider?
- 4.6: Getting the most out of AI feedback
- 4.7: AI for research assistance

This page titled [4: AI Strategies](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative) .

4.1: How do we tell uses of AI that help learning from ones that hurt?

Everyone knows chatbots can be used to cheat. They can also be used to support learning. Sometimes it's easy to tell the difference: if AI does all the work and we do nothing, we don't learn. If we get AI to make up hard questions, and we have to come up with answers, we do learn. But often it's not that clear.

There are so many forms of AI assistance. In her [Vox.com](#) video [AI can do your homework. Now what? Students and teachers grapple with the rise of the chatbots](#), journalist Joss Fong gives us a taste of myriad possibilities:

- Answers to a homework question.
- Background information on a topic.
- Definitions or explanations of a concept.
- Sources to find more information.
- Summaries of readings and lectures.
- Study guides for an exam.
- Ideas for how to respond to an assignment.
- Instructions for solving a problem.
- An outline for a paper or presentation.
- Examples, analogies, and counterarguments.
- A draft of a paper or a discussion post.
- A script for a presentation.
- Feedback on their work.
- A revision of a text to improve it.
- A revision of a text to change its word count

So what's a good idea, and what's not? Here, I want to suggest one principle and one rule of thumb that could guide us in the maze of options.

A principle: focus on how AI use affects your thinking

As Fong points out in the video, it can feel like the point of school is to produce work that meets criteria. Really, of course, the point is what happens in our brains when we do the work. When I develop guidelines for AI use, I try to judge how each possible application might help or hurt the learning goals of the assignment. I'm more interested in mental processes than in abstract ideas about what constitutes cheating.

Self-monitoring learning

Figuring out whether a particular use of AI helps us learn often requires a lot of self-awareness. When it's not obvious whether it's helping more than it's hurting, you'll need to act like something of a scientist: experiment, reflect, and adjust your practices. Ask yourself,

- Did using AI in that way help me get through a stuck place?
- What other strategies could I have tried?
- Did the AI use make me feel more or less engaged?
- What could I do again or do differently next time?

Of course, your teacher may be able to guide you. (See [How do I know which use of AI is allowed and which isn't?](#)) But they may not have thought of every use, and they may leave some decisions to you.

The value of friction

The tricky thing about monitoring our learning is that we tend to underestimate how much we're learning when we're frustrated. Joss Fong discusses a study by Deslaurier et al. where students who listened to a polished lecture rated their learning high but actually did worse on the test. Students who engaged in active learning with the same teacher were frustrated and ranked their learning low but did better on the test.

As Fong observes, “Whenever we try and judge if a learning experience is productive or not, the strongest metacognitive cue that we use is perception of fluency. Fluency is when information is going down easy. It’s well presented, it’s organized, it’s convenient.” But, as she explains, “Education researchers have this term ‘desirable difficulties,’ which describes this kind of effortful participation that really works but also kind of hurts. And the risk with AI is that we might not preserve that effort, especially because we already tend to misinterpret a little bit of struggling as a signal that we’re not learning.”

So it’s worth remembering that a frustrating path might be more fruitful than an easier AI-assisted path. Fong suggests using AI in ways that allow us to do harder work. Rather than using “a chatbot to avoid reading a challenging text... you could use it to work through that text and help you get more out of it.” We don’t need to be perfect at this; we just need to pay attention and keep trying. Gradually, we’ll get better at self-awareness and strategy.

A rule of thumb: use AI for tutoring-style assistance

To reduce the complexity of these decisions, it helps to have a rule of thumb to turn to. If you’re considering a specific use of AI, ask yourself, would an ethical tutor agree to do this?

Generally speaking, tutors expect that student work should be student ideas and student words. They don’t complete part of the assignment, but they do give feedback, examples, and explanations. Tutors get training to help them make the call on gray areas: is it okay to suggest a new outline for your essay or should they ask you questions to help you come up with one on your own?

When you go to a tutor, they decide what’s helping and what’s overhelping. With a chatbot, you have the burden and freedom of drawing that line. You might not know exactly where to draw it, but you probably have some sense of what a human tutor will and won’t do. If you can’t see a human tutor rewriting your sentences to make them sound sophisticated, it’s probably not a good idea to ask AI to rewrite them.

For more guidance on using AI for tutor-style alternate explanations, examples, quiz questions, or feedback, see [AI for Tutoring-Style Assistance](#). (And remember, this tutor is unreliable and possibly biased.)

AI for task completion? Not unless the teacher has designed for it

What about using AI to complete part or all of an assignment? For example, what about asking AI to write a first draft? Some argue that since we’ll likely use AI in future workplaces, we should use it to do homework. However, as we’ve noted, the purpose of schoolwork is to change our brains. (There’s no other great reason to add to the pile of completed assignments in the world.) The purpose of work is to get things done for someone who pays us.

Some teachers design assignments to invite AI collaboration. They identify learning goals students can achieve by prompting AI and responding to what it gives. Then they invite students to use AI to complete some part of the assignment. Students check, revise, and add their own words and ideas. Ultimately, they are responsible for the quality of the result.

If your teacher has not explicitly designed for this, however, then using AI to actually do what the teacher assigned will likely get in the way of learning. Ask yourself,

- What is the assignment supposed to teach?
- What might I learn if I do what I was going to ask the AI to do?

It’s safer to stick with tutoring-style assistance.

Further reading

- [AI can do your homework. Now what? Students and teachers grapple with the rise of the chatbots](#), produced by Joss Fong for [Vox.com](#)
- [Measuring actual learning versus feeling of learning in response to being actively engaged in the classroom](#) by Louis Deslauriers, Logan S. McCarty, Kelly Miller, and Greg Kestin
- [AI Assessment Scale](#) by Mike Perkins, Leon Furze, Jasper Roe, and Jason MacVaugh
- [Student Use Cases for AI](#) by Ethan Mollick and Lilach Mollick
- [AI Study Buddy](#) from Northwestern University
- [Principles for Using AI in the Workplace and Classroom](#) by Joel Gladd



Login with LibreOne to view this question

NOTE: If you typically access ADAPT assignments through an LMS like Canvas, you should open this page there.

Login

This page titled [4.1: How do we tell uses of AI that help learning from ones that hurt?](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) ([ASCCC Open Educational Resources Initiative](#)) .

4.2: AI for tutoring-style assistance

Before you ask AI for assistance, try asking yourself “Could a human tutor help me in this way without doing my work for me?” Most of us recognize that human tutors can do all kinds of things to stimulate our learning, and they also have to draw a line between helping and overhelping. If we use that line as a standard, we can be pretty sure AI will not interfere with our learning.

I do want to make clear that I’m not saying AI assistance replaces human tutoring or is better than human tutoring. Human tutoring offers actual human connection, empathy, encouragement, witnessing and other forms of support that have been shown to have a huge impact on learning. Your human tutor might actually be interested in what you have to say and might care about you as a person and a learner; AI can say nice things, but it isn’t capable of those feelings or relationships.

That said we don’t always have access to human tutors. And not every human tutor is a good fit for every student. They might not be as flexible or familiar with particular approaches to learning assistance that you want to try. These are reasons why I see AI tutoring-style assistance as a valuable supplement. In my writing classes, I assign students to visit a human tutor first and later reflect on AI feedback*.

Get the chatbot to explain differently

Sometimes we turn to tutors when we just don’t get it, when the course materials don’t make sense. The tutor explains in a different way and keeps trying different explanations until we get it. We can ask a chatbot for that. And with a chatbot, we can get ultra specific about what flavor of explanation we want.

For example, we might ask for an explanation that fits our intellectual strengths or how smart we’re feeling that day:

- “Explain it for a five-year-old/a twelve-year-old/a graduate student.”
- “Explain it for a physics major who doesn’t get poetry,”
- “Explain it for an English major who doesn’t get math.”

We can also ask for engaging explanations when we’re having trouble related to the material:

- **Why it matters:** Ask it to help you understand the relevance or application of the concept to help you stay motivated to learn. or example, if you were studying the theory of mind in preschool-aged children, in addition to asking for explanations of what that is, you could ask “Why does it matter if a kid has a theory of mind or not?”
- **Style or genre:** Ask it to explain it in a creative way that’s more fun or that reaches you differently, like a podcast (try NotebookLM), a spoken word poem, a rap, a sonnet, an image, a short story, or a game (This would be tough for a human tutor to produce on the spot.)
- **Examples:** Ask it to tailor examples to things you care about or are interested in. Maybe you are learning about opportunity cost in an economics class and you ask for an example related to sustainable fashion.
- **Comparisons or metaphors:** For a concept that doesn’t have easy-to-relate-to examples, you can ask for a comparison or metaphor instead. Let’s say you are learning about covalent bonding in chemistry and you ask the chatbot to make up a metaphorical explanation involving dogs while still conveying the essential concepts. (Here’s [a ChatGPT answer involving dogs sharing bones](#) instead of atoms sharing electrons.)

Get the chatbot to ask you questions and help you practice

You can also use a chatbot to test and expand your understanding of the course material, such as before a test, before class discussion, or before a meeting with an instructor. Tell it to make up sample questions based on your topic and course materials. If you’re preparing for a test, give the chatbot any information the instructor has given you about the test.

For example, here’s [a sample chat where I asked it to ask me questions about logical fallacies](#) based on a textbook chapter (the chapter was published online under an open license, so I wasn’t violating any rights by uploading it to the chatbot).

If you have to do an oral exam or you need to interview someone for an assignment, you can ask a chatbot to simulate the exam or interview to help you prepare. The AI Pedagogy Project at Harvard’s metaLAB gives an example of how students preparing to interview heritage speakers of Spanish could [practice their interviewing skills with a chatbot](#) before the real interview. (Consider using voice mode if you prefer to do this out loud)

Learn by teaching the chatbot

They say if you want to learn something deeply, teach it. What if you pretend to be the teacher while the chatbot pretends to be the student? It explains something to you (maybe imperfectly), and then you tell it how to improve the explanation. In [“AI as Learner: Challenging Students to Teach,”](#) Ethan and Lillach Mollick provide a prompt for this. They share a [chat session where the student critiques the chatbot’s explanation of “first-mover advantage”](#) in business.

How do you check AI “tutoring”?

Wait a minute, though. Given all the [cautions](#) we’ve heard about bias and inaccuracy in chatbots, we’re always being told to check what comes out of them. But how the person who needs a tutor supposed to know whether the tutor is right? As Leslie Allison says in [“AI Can Do Your Homework, Now What?”](#) “The less you know about something, the more likely you are to be convinced by ChatGPT’s answer.”

There’s no perfect solution to this challenge. It might help to think of a crazy cousin who knows a lot but gets mixed up. They’ll often give accurate explanations and examples, but sometimes they will sound great and be wrong in weird ways. Would you still turn to them for tutoring assistance? How would you know what to trust?

The good news is that there are some ways you can find out if it’s wrong. It takes a little time, but there are faster ways to do it, and the time is not wasted; it’s probably helping you learn the same concepts you were trying to get help with.

Ways to investigate accuracy

Here are some strategies to help you assess AI tutoring-style assistance.

- **Does it make sense to you?** Trust your own thinking and investigate when something doesn’t seem quite right.
- **Does it match what’s in your course materials?** Circle back to readings, notes, and textbook explanations. You can ask the chatbot to double check against your course materials if those are publicly available and quote from them to show how they support its explanation or example.
- **Let the chatbot critique itself.** You can ask AI to play the role of an expert and double-check itself. For example, when [I asked ChatGPT to explain covalent bonding](#), it compared atoms sharing electrons to dogs sharing bones. I asked, “How might an expert chemist critique this explanation? In what way is it misleading?”
- **Check against credible web sources (with chatbot help).** Search, or ask the chatbot to suggest a link to a credible website that supports its claim or explanation.
- **What does a teacher, human tutor, or classmate say?** Keep a human in the loop. AI assistance and human discussion can complement each other.

A personal example

Here’s an example of a way I used it in my own learning. “Embedding” is a term for something key in the structure of large language models. I’ve heard “embedding” for years, seen it explained, but not really understood. Recently, I read it somewhere and got frustrated. I knew it has something to do with the math that language models use to represent and analyze text. But what was embedded in what? I felt intimidated and didn’t want to invest the time to absorb a technical article. Besides, I was skeptical that such an article would explain it in the simplest way possible.

So I [asked ChatGPT to explain embeddings in five ways](#) that a twelve-year-old could understand. The results appealed to me and reminded me that I had read that embeddings were vectors, ordered sets of numbers that represented words. Here are two

- “Imagine every word took a huge personality quiz with hundreds of questions like “Are you positive or negative?”, “Do you go with animals?”, or “Are you something people can touch?” The answers to all those questions are stored as a list of numbers—that’s the word’s embedding.”
- “Embeddings are like a secret number code that the computer uses to remember what a word means instead of how it’s spelled. So instead of ‘dog,’ the computer sees something like [0.2, -0.1, 0.7, ...]. That code helps the computer figure out how “dog” relates to ‘puppy’ or ‘cat.’”

Then, when I was feeling a little more confident, I asked it for an [explanation aimed at an educated layperson with no background in computer science](#). The explanation made some sense. It ended with the evocative lines, “Embeddings are how a language model

turns words into math so it can work with meaning. Think of them as coordinates on a map of ideas that help the model "navigate" language.” ’

Next, I [asked it how a computer scientist would critique its explanation](#). While I didn’t understand everything in the critique, I did appreciate some clarifications that made sense given credible pieces I had read on LLMs. For example, it explained that the same word could have a different embedding, or sequence of numbers, depending on context, which made sense because a word can mean different things depending on context.

At the end of this chat session, I didn’t feel like I had mastered embeddings, but I did feel better equipped to read and respond to human discussions about them. I was intrigued to learn more.

4.2: AI for tutoring-style assistance is shared under a [not declared](#) license and was authored, remixed, and/or curated by LibreTexts.

4.3: Reasons not to let AI brainstorm, write, revise, or edit (in college)

Most people struggle with writing, and most people wish there were an easier or faster way to write. It can be very tempting to turn some of the work over to AI. It can seem like AI is leveling the playing field by giving students who have a hard time a chance to boost their grades. But what seems to help might hurt in the long run.

Often, writing is hard because doing it involves hard thinking, and that is where we learn. (See [AI can generate decent-sounding text. Do we still need to learn to write?](#))

Some writing teachers I respect are experimenting with allowing AI assistance with brainstorms, drafts, and revisions, resulting in hybrid, part-AI writing. The [Elon University Student Guide to AI](#) suggests, “Think of the writing process as a human-AI loop. You remain in charge of the writing but use AI at certain stages of the process, working back and forth to create the final paper.”

Certainly, some learning and thinking can happen if we let AI come up with ideas and words and then add or make changes. AI outputs can stimulate thinking as well as replacing it. But in my view, on balance, letting AI do some of the writing takes away chances for students to develop voice, confidence, and critical thinking.

Brainstorming, organizing, writing, and revising in college are essential. Yes, “AI is here,” but that doesn’t mean we have to use it for everything. Below, I’ve explained why I ask students to use AI in the writing process only for feedback that leaves it to them to make changes to their drafts.

Note: I leave it to learning disability specialists to identify cases where specific uses of generative AI in the writing process should be considered accommodations. For students with writing disabilities, some kinds of friction in the writing process may be more destructive and limiting than useful and necessary.

Thinking by brainstorming and outlining

In brainstorming and outlining, we wrestle with ideas. If AI supplies ideas and organizes them, we miss chances to develop mental muscle and intellectual confidence.

Imagine you are starting on a paper about ways to prevent teen social media addiction. You’ll rack your brain for examples you’ve heard of and approaches you’re curious about. Maybe you do some research and find studies on the effectiveness of media literacy programs versus apps that limit access to TikTok and other platforms. To get some words down, you’ll have to reflect on what stood out to you in the research and how it connects to your prior assumptions and examples you’ve seen. You might get stuck, but there are plenty of strategies that help, from voice typing to timed writing with partners to giving yourself permission to “write the worst junk in the world,” as Anne Lamott puts it. (See the section on [Brainstorming](#) in my textbook, *How Arguments Work: A Guide to Writing and Analyzing Texts in College* for more strategies.)

Once you have a mess of a brainstorm, if you try to make an outline or organize it into paragraphs, you’ll need to reread what you wrote and notice when you switched from one point to another. You’ll reflect on which points seem important. Maybe you’ll consider which example from the beginning of the brainstorm relates to the study you mention later and what point they both illustrate that could become a topic sentence. As you sort these things out, you start feeling like you have a map in mind and you know much better what you want to say.

Chatbot rewrites can make a piece of writing worse and kill your voice

If you ask a chatbot to polish your essay, it might just butcher it. AI rewrites can sound academic and sophisticated, but AI often changes the meaning as it edits. It can be harder when you have a polished sounding rewrite in front of you to figure out whether it’s what you want to say or how you would say it. For this reason, I don’t ask AI for rewrites myself.

Here’s an example where a ChatGPT rewrite makes student writing worse. I gave ChatGPT the introductory paragraph from “We Don’t Care About Child Slaves” by Jennifer Vo-Nguyen of Portland State University, a student essay selected for inclusion in a college writing textbook, [EmpoWORD: A Student-Centered Anthology and Handbook for College Writers](#).

Here is the original student-written opening from [EmpoWORD: A Student-Centered Anthology and Handbook for College Writers](#):

“When you walk into the mall or any department store, your main goal is to snatch a deal, right? You scout for the prettiest dress with the lowest price or the best fitting jeans with the biggest discount. And once you find it, you go to the checkout and purchase it

right away. Congratulations—now it’s all yours! But here’s the thing: the item that you just purchased could have possibly been made from the sweat, blood, and tears of a six-year-old child in Vietnam. According to the International Labour Organization (ILO), one in ten Vietnamese children aged 5 to 17 are slave workers, and Vietnam is the second biggest source of imported goods to the US.”

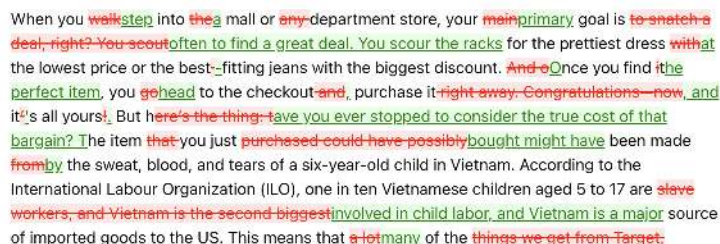
Here is the ChatGPT rewrite (See also the [full chat session transcript](#)):

“When you step into a mall or department store, your primary goal is often to find a great deal. You scour the racks for the prettiest dress at the lowest price or the best-fitting jeans with the biggest discount. Once you find the perfect item, you head to the checkout, purchase it, and it’s all yours. But have you ever stopped to consider the true cost of that bargain? The item you just bought might have been made by the sweat, blood, and tears of a six-year-old child in Vietnam. According to the International Labour Organization (ILO), one in ten Vietnamese children aged 5 to 17 are involved in child labor, and Vietnam is a major source of imported goods to the US.”

The rewrite sounds fine, even a little smoother than the original. But it also sounds less alive. It wouldn’t get a better grade in my class. Most teachers hope students will write something that expresses their voice—their curiosity, interests, strong opinions, passions, or feelings of uncertainty. That’s where the energy is—that’s what we like to read!

ChatGPT’s rewrite also detracts from the student’s point. Jennifer Vo-Nguyen clearly intends to protest what she sees as slavery, not just labor. Using the word “slave” expresses moral outrage more than “child labor.” “Slavery” echoes the language of her sources, such as a BBC article titled “Vietnam’s Lost Children in Labyrinth of Slave Labour.”

Below you see ChatGPT’s edits in the screenshot below with the EditGPT extension.



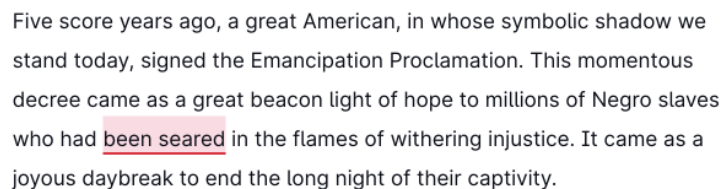
When you ~~walk~~^{step} into ~~the~~^a mall or ~~any~~^{any} department store, your ~~main~~^{primary} goal is ~~to snatch a deal, right?~~^{to find a great deal}. You ~~scout~~^{scour} ~~often~~^{the racks} for the prettiest dress ~~with~~^{at} the lowest price or the best-fitting jeans with the biggest discount. ~~And~~^{Once} you find ~~the perfect item~~^{it}, you ~~go~~^{head} to the checkout ~~and~~^{and} purchase it ~~right away~~^{right away}. ~~Congratulations—now, and it's all yours!~~^{Congratulations—now, and it's all yours!} But ~~here's the thing: have you ever stopped to consider the true cost of that bargain?~~^{here's the thing: have you ever stopped to consider the true cost of that bargain?} The item ~~that~~^{that} you just ~~purchased could have possibly bought~~^{bought} might have been made ~~from~~^{from} by the sweat, blood, and tears of a six-year-old child in Vietnam. According to the International Labour Organization (ILO), one in ten Vietnamese children aged 5 to 17 are ~~slave workers, and Vietnam is the second biggest~~^{involved in child labor, and Vietnam is a major} source of imported goods to the US. This means that ~~a lot many~~^{a lot many} of the ~~things we get from Target~~^{things we get from Target}.

The red crossouts seem so authoritative. What if this student had taken the advice because she didn’t feel confident? Would the teacher still have chosen the more boring, AI-polished version as a model for inclusion in a textbook?

Grammar suggestions can get in the way, too

Grammarly’s basic identifications of grammar errors and suggested fixes can be helpful (I use them myself), but it’s key to stay skeptical. Sometimes the grammar suggestions are wrong, and sometimes they change the meaning completely. Remember that Grammarly is always trying to sell you its product by suggesting that you need it because it knows better.

Grammarly even claims it can do better than Dr. Martin Luther King, Junior. His [“I Have a Dream” speech](#), is surely one of the most powerful and influential speeches in American history, but when I put it into Grammarly, the app claimed to find 16 errors and 5 ways to “improve your text” for clarity. It dangled the promise of 18 “Pro tips” for paying subscribers. Some of its suggestions, such as suggestions for comma insertion, were fine but not important. Others were just wrong. Here in the first paragraph, it highlighted the phrase “been seared” in red:



Five score years ago, a great American, in whose symbolic shadow we stand today, signed the Emancipation Proclamation. This momentous decree came as a great beacon light of hope to millions of Negro slaves who had **been seared** in the flames of withering injustice. It came as a joyous daybreak to end the long night of their captivity.

Below is a screenshot of the Grammarly explanation. It suggests that MLK should “Correct the verb” by changing “been seared” to “suffered.”

...slaves who had ~~been-seared~~ **suffered** in the flames...

Accept

Dismiss ...

It's true that "been seared" is passive voice, and sometimes passive voice is overused or less effective than active voice. That might be why Grammarly flagged it. Here, however, passive voice is needed. "Seared" vividly evokes suffering through burning, connecting to the "flames" mentioned later in the sentence and adding to the drama and specificity of the sentence. It reminds us of "searing," which describes intensely damaging and fiery effects, as in "a searing critique." ChatGPT's "suffered" is more generic. Given the powerful, epic, biblical style Dr. King was aiming for, "seared" fits beautifully and should not be sacrificed to the Grammarly "fix."

Chatbot sentence-level style and grammar suggestions can also get in the way of more interesting, varied, and authentic expression. Chatbots are biased toward standardized English (See [Don't trust AI: it's biased](#)). I tested this by giving ChatGPT the opening paragraph of Dr. Vershawn Ashanti Young's famous essay "[Should Writers Use Their Own English?](#)" When I instructed it simply to "Revise this," it [rewrote the paragraph](#) without any of the "Black English" (as Young calls it there) that Young had so consciously and defiantly embraced. When I gave it the same passage and asked "How could this be improved?" it again rewrote the passage without any Black English. This time it also [injected an admonition](#) to ensure "proper grammar."

Removing linguistic variation goes against the idea of "[Students' Right to Their Own Language](#)," a concept affirmed by College Composition and Communication (1974) of the National Council of Teachers of English in 1974 and [reaffirmed in 2003 and 2014](#). They declared, "We affirm the students' right to their own patterns and varieties of language -- the dialects of their nurture or whatever dialects in which they find their own identity and style. . . ."

As we'll see in [Getting the most out of AI feedback](#), some uses of AI feedback may support linguistic justice and help you build confidence in your voice and the way you might choose to reach readers through language variation and code meshing. But we've got to keep a tight rein on AI to make that work.

Further reading

- [The introduction to *How Arguments Work: A Guide to Writing and Analyzing Texts in College*](#), including [Why Study Argument?](#) and [A Closer Look at Fast and Slow Thinking](#)
- [The Seductions of AI for the Writer's Mind](#) by Meghan O'Rourke in *The New York Times*, July 18, 2025

This page titled [4.3: Reasons not to let AI brainstorm, write, revise, or edit \(in college\)](#) is shared under a [CC BY-NC](#) license and was authored, remixed, and/or curated by [Anna Mills](#) ([ASCCC Open Educational Resources Initiative](#)) .

4.4: Getting the most out of AI (prompting)

A chatbot is such an open-ended thing. The possibilities are endless; you can use any words you want to ask it for help. So where do you start? What are the best techniques? And how long will it take to learn them? Will it be too difficult if we're not coders?

Here's some good news: in my opinion, there's no special technical knowledge needed to use chatbots well. Even though some are referring to prompting skills as "prompt engineering," there is no evidence that "engineering" of any kind needs to be involved. Reading, writing, and critical thinking are the main skills needed.

You may see complicated formulas with acronyms for the kinds of things to include in your prompts, and there's nothing wrong with those. But they're absolutely not required. Very simple prompts that just say what you want in the words you would naturally use may serve you just as well.

I've come to this conclusion through my own practice, bolstered by the writing of Ethan Mollick, a well-known voice on AI in education and business, author of the New York Times bestseller *Cointelligence: Living and Working with AI* and the newsletter [One Useful Thing](#). He reviews current research into prompt engineering and spends countless hours experimenting with the latest models. Mollick advises, "Don't aim for perfection - just start somewhere and learn as you go" ("[Getting started with AI: Good enough prompting](#)").

Be specific

Think of more ways to describe what you are looking for. How can you be more specific? Below are some examples of ways you can get more specific with a chatbot. Again, none of these are required, so don't let this list slow you down as you prompt!

- **Context:** What's the project you are working on? For example, you could specify that you are preparing for a test on X or developing a scholarship application for a nursing program. You can copy and paste from any description of the test or application. The context might include the **genre** you are writing in, like a business memo, a medical research paper, or a political speech.
- **Role:** Is there a **role** you want the chatbot to play? If you could choose any kind of human to help you, what kind of person would you choose? You can ask the chatbot to (try) to play that role. Some possibilities: expert tutor, study skills coach, scientist, editor, curious skeptic, devil's advocate. Notice that some of these describe a professional role (tutor) and some just describe an attitude (curious).
- **Style:** Is there a particular style you prefer or a style that's expected? For example, when I was working on my professional bio, I told Claude, "I don't want it to sound like I am full of hot air. I don't like professional bios that seem generic and exaggerated like 'innovative thought leader.'" You can describe the style by giving the name of a writer you want it to imitate or by throwing in adjectives like "academic," "conversational," or "simple but elegant." One resource you can read and even attach to your chat session is the chapter [Style: Shaping Our Sentences](#) from my free and open textbook *How Arguments Work: A Guide to Writing and Analyzing Texts in College*.
- **Tone:** Style and tone overlap, but tone often refers to the emotional quality you are going for. For example, "optimistic," "earnest," "reverent," "scathing," or "melancholy." For more, see [Tone](#) from my free and open textbook *How Arguments Work*.
- **Examples:** Just as you might with a person, give the chatbot an example of the kind of thing you are looking for. You can paste in an extended sample if you have one, or you can describe the example briefly or allude to something well-known. For example, you might say "I'm looking for constructive feedback on... Here's an example of what I consider helpful feedback..."

Ask it to ask you questions about what you want

Sometimes we don't have a clear idea of what we want or it feels overwhelming to try to articulate it. In those cases, one approach is to ask the chatbot to ask us questions about what we're looking for. For example, when I wanted a sample argument to give my students to analyze for their final exam, I said "I need to create three new sample arguments with many of the same qualities as the samples below but with different themes. I would like two of the arguments to be about the future of AI in writing instruction. Ask me questions to help you develop the sample arguments."

Get AI to take it step by step

Some have called this "chain of thought prompting." (It's kind of fun to have a technical term, but remember, people are just experimenting with these bots and making up catchphrases.) This can be as simple as literally telling the chatbot: "take it step by

step.” Or you may get better results by telling the chatbot what the steps are. What steps would a human need to take to get to the answer? For example, if you want ideas on how you might responding to a letter from your health insurance company, you might give it the following chain of instructions:

1. Start by asking the chatbot to summarize main points from the letter.
2. Then tell it to consider what kinds of further information it needs,
3. Ask it to browse the internet to find credible sources that supply that information
4. Ask it to analyze those sources and summarize how they help us interpret the letter and respond to it.
5. Finally, ask it for a list of five possible strategic responses.

Prompting the system step by step provides you an opportunity to catch problems in a chatbot’s approach and correct them as you go. In the above example, if it misrepresents a main point from the letter, you want to catch that before you let it go on to come up with strategies for responding.

Push the AI to do better

Don’t judge a chatbot by its first response. Chatbots won’t tell you the same thing every time. And they don’t get their feelings hurt. So if you don’t like something about what the chatbot gives you, good! Tell it. It may be able to give you something better.

This is so simple, but somehow it isn’t intuitive for most people when we first interact with a bot. We might be tempted to turn away and later tell a friend “it got X wrong” or “it’s style is too bland” or “its answers are simplistic.” Instead, we could tell it, “That’s wrong because... Try again.” Or “Give me another version that is more nuanced.”

If you can’t put your finger on what’s wrong, then you have some key information that might help you get a better result. Just tell it what’s wrong. Ask it for another version that corrects the problem.

Even if you are not sure what is bothering you about what it gave you, if you are underwhelmed, you can ask it to give you another, more insightful response. Sometimes I just say “Please try again.”

Lean in to your reading and writing skills and keep building expertise

Many people are wondering how best to prepare for future workplace uses of AI. One thing we can be sure of is that you need basic understanding of your field and critical thinking skills to work with AI. To prompt well, you need to know what you are looking for and what sounds good but isn’t right. To know that, you need subject matter expertise.

That means what you’ve already been learning in school and what teachers already know how to teach will help you use AI. Writing and reading practice and understanding of a field will all help you get more out of AI. So will awareness of your own learning and thinking processes.

Here’s another way to think of it: the things you’re learning without AI will help you to use AI later on. All of your studies prepare you for capable prompting and iteration down the line.

Be playful and curious

Working with a chatbot is so open-ended—you’re in the sandbox, and the possibilities are endless. There’s not one way to proceed. One way to think about it is as a play space. When you’re stuck on a task or on how to get help, enjoy the freedom you have to try lots of different things, complain about what the AI is giving you, boss it around, try creative or strange prompts that strike you intuitively.

We can learn a lot about language models by trying weird things with them. For example, Ethan Mollick gave the Claude chatbot the text of a book and [asked the bot to “remove the squid.”](#) In another experiment, he [kept prompting it with the phrase “garlic bread.”](#) The results are entertaining, but the chat sessions also show how we can think outside the box when prompting, how the language model tries to find an existing pattern to match these odd prompts, and how inconsistent its responses are.

A playful approach to prompting is not just more fun; it means flexibility, curiosity, and openness, which are likely to lead to better results.

Further reading

- For a basic introduction with screenshots, see [How to Prompt AI Chatbots](#) by Joel Gladd and Liza Long
- For a quick list of prompting tips, see [Getting started with prompts for text-based Generative AI tools](#) from Harvard University Information Technology
- For a substantive, friendly discussion of what to try and what not to worry about, see [Getting started with AI: Good enough prompting](#) by Ethan Mollick
- For an approach to prompting that involves developing an imagined “relationship” with AI, see Ethan Mollick’s [On the necessity of a sin: Why treating AI like a person is the future](#)
- For a more technical look at prompt engineering, see [Anthropic’s prompt engineering overview](#) and [prompt library](#)
- For prompt engineering mini-courses of one hour or more, see [LearnPrompting.org](#)
- For a playful, humorous approach to AI systems, see Janelle Shane’s [AI Weirdness blog](#) and her book [You Look Like a Thing and I Love You](#).

This page titled [4.4: Getting the most out of AI \(prompting\)](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative) .

4.5: Which AI tools should I consider?

I find that a lot of people who haven't used AI much feel some anxiety about figuring out what to use. Choosing can be pretty simple. You don't have to read this whole article if you just want a starting place; you can just read the beginning and stick with the primary chatbots. Beyond those, I've included informal recommendations based on what I use and what I can imagine being useful for students.

Remember all the [cautions](#) around privacy, bias, energy use, and more described earlier as you consider whether and how to use these systems.

Try at least one of the three main chatbots

There are a lot of apps out there, but they are unlikely to give better results than the chatbots from the three leading AI companies. So don't assume you're missing out if your friend is using an app you've never heard of that keeps advertising to you on YouTube. Many, perhaps most other AI apps use the same underlying technology from those same companies but with a different focus or user interface.


Since we can talk to these systems about what we want from them, general chatbots can often produce good results for many different kinds of tasks; you won't need a specialized app. In the process of using a general chatbot, you'll get a lot of practice with prompting.

The three most powerful chatbots have similar user interfaces and accessible free versions:

- [Claude](#) from Anthropic. I have heard many highly educated people and people deep in the AI world express a slight preference for Claude when it comes to writing, and that matches my experience so far. I've also heard of others who prefer ChatGPT or Gemini.
- [ChatGPT](#) from OpenAI.
- [Gemini](#) from Google.

Choose which model your chatbot runs

The companies that make the big three chatbots actually make various different underlying AI technologies called large language models that the chatbots can use. The quality you get from a chatbot completely depends on which model it is running. To save money, companies don't often set chatbots to default to the most sophisticated underlying models. So take a moment to see which model it's using. Often, you can select a more sophisticated model without upgrading your account. For example, in the screenshot below, the user has selected o3 instead of the default GPT 4o. OpenAI's o3 and Claude Opus 4 are examples of so-called "[reasoning models](#)," a newer type that has undergone extra training for complex reasoning tasks and also spends more time trying to work through a request step by step.



Screenshot of the ChatGPT model selection interface as of July 2, 2025

Other AI apps you may want to explore

I use the systems below in my own work at times. They are (mainly) based on the same underlying models from Google, OpenAI, and Anthropic, but offer a different user interface.

- [Bartman](#) is a chatbot “trained with diverse histories & inclusive voices” especially black and brown people. It is designed around search chat combination. Some love the interface and use experience. It is good for focusing on finding and synthesizing information drawn from searches of YouTube and Reddit. It also has practice test and flash card features.
 - [NotebookLM](#): Google’s interface designed for writers allows you to upload more files as context for your chats and will autogenerate fairly engaging podcasts based on those files.
 - [Copilot](#): Microsoft’s chatbot is similar to ChatGPT, and your school may offer access to it with more data protections than you would get with a personal ChatGPT account. However, I have heard several people suggest it’s not of the same quality as the big three chatbots.
- [AI apps for writing or studying with voice and audio](#)

We think differently when we listen and talk out loud than when we read and write. You might want to experiment with interacting with AI through audio or with using AI to help you switch between talking and writing.

ChatGPT and Gemini are said to have the best voice modes; Claude has just added voice mode as well. I sometimes talk aloud with ChatGPT while exercising or doing chores. I find it useful as a thought partner that can browse the internet and help me explore an idea while I’m moving my body instead of sitting in front of a screen.

If you want to listen to a text instead of reading it, the Microsoft Edge browser has a pretty good built-in [Read Aloud feature](#) that lets you listen to web content in a variety of synthetic voices.

- To dictate my writing, I use the following:
- [Voice typing](#) in Google Docs
- [Notes in iPhone](#) allows dictation if you select the microphone icon at the bottom of the keyboard.
- [Otter AI](#) is designed for meetings, but I have used the free version for personal dictation despite a somewhat confusing and complex interface.

If I’m dictating into Notes or Otter, I will often copy the transcript into a chatbot and say “Please format this transcript and clean up errors without changing the wording.”

AI for academic research assistance

See also [Pros and cons of AI for research assistance](#).

- [Stanford’s STORM](#) generates Wikipedia-like articles on a topic of your choice. The Co-STORM version integrates multiple perspectives.
- [Elicit](#) and [SciSpace](#) help you find, summarize, compare and analyze academic research.
- [Consensus](#) also helps you navigate academic research with a focus on assessing the level of agreement among scholars on your question of interest.
- [Keenious](#) helps you find more academic research papers related to a given paper.
- [Scite](#) lets you “see how a publication has been cited by providing the context of the citation and a classification describing whether it provides supporting or contrasting evidence for the cited claim.”
- [MoxieLearn](#) offers access to AI platforms, coaching, and prompts to support academic research and writing. It is a bit more costly than others but may be worth the investment for graduate students.
- [Undermind](#) “highlights the precise papers you should focus on and gives a clear explanation for each decision.”
- [ConnectedPapers](#) lets you explore relationships between papers in a graph.
- [ResearchRabbit](#), a “Spotify for research,” includes literature review and personalized alerts when new research is published in your area.

Custom chatbots for education

A custom chatbot is just a chatbot with an extra prompt behind the scenes shaping how it responds. Here are just a few that I use or have developed.

- [Deep Background: Fact-Checks and Context](#) gives you an analysis of any claim you put in with links to sources. A custom bot on the ChatGPT platform by Mike Caulfield, creator of the influential SIFT model for verifying claims online.
- [Contradictory chatbot](#) gives answers much like a generic chatbot except that it is instructed to always provide three reasonable but incompatible answers to each query. The purpose is to encourage users to practice treating plausible AI outputs with skepticism.
- [Contradictory chatbot for research](#) will browse the internet and give three contradictory but reasonable answers, each with a supporting source link.
- [PlayLab apps for college students created by educators](#) (filter by subject matter). PlayLab is an educational nonprofit.

Commonly recommended AI apps for education

I don't have direct experience with the systems below, but I know people who've had good experiences.

- [Boodlebox](#) is a chat platform that uses the best AI models with prompt guidance for education. Individual students can use it. Many educators I know on social media are getting excited about this platform; I haven't experimented.
- [Khanmigo for Learners](#) from the nonprofit Khan Academy is a popular AI tutoring platform.

Other AI apps you may wonder about

I don't have direct experience with the systems below, and I haven't heard the educators and AI experts I follow on social media recommending them.

- [GrammarlyGo](#): The Grammarly app and extension include capabilities like generating and revising text.
- [Grok](#) from X
- [Llama](#) from Meta
- [DeepSeek](#) from Hangzhou DeepSeek Artificial Intelligence Basic Technology Research Company
- [ManusAI](#), an agentic system designed to browse and take actions online on the user's behalf. (I'm not going near this for the moment at least—think of what could go wrong.)

Further resources

- [Using AI Right Now: A Quick Guide: Which AIs to use, and how to use them](#) by Ethan Mollick
- [AI Models and Tools](#), an extensive, frequently updated list with a focus on education from Jose Antonio Bowen

This page titled [4.5: Which AI tools should I consider?](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by Anna Mills ([ASCCC Open Educational Resources Initiative](#)) .

4.6: Getting the most out of AI feedback

Why AI feedback?

As I see it, reflecting on AI feedback on your work is a relatively straightforward and low-risk way to use AI. It allows you to get help and to build your prompting skills without missing out on thinking and learning. It's a powerful alternative to letting AI write or rewrite for you.

By AI feedback, I mean comments that stimulate your thinking as you revise. I don't mean chatbot rewrites or chatbots feeding you ideas or sources. A tutor shouldn't tell you what to say, and a chatbot shouldn't either. But both can support you to figure out what changes you want to make.

Reflecting on AI feedback means you can stay in control, keep your voice, and claim your words and ideas. And it can help you build confidence in your own judgment if you stay skeptical of the suggestions. Some AI advice will be bad, or it will be fine but not for you. You are the one who knows what you want to communicate to your readers. If you develop a habit of questioning AI feedback, you'll be more likely to question AI in other contexts, which will serve you well in the workplace and in your personal life. People who can see where AI is wrong can improve on it or set it aside when something else is needed.

Another benefit to engaging with AI feedback is that you are unlikely to violate an academic integrity policy. If you use AI for ideas or drafting or rewriting, there will be pitfalls and gray areas around what is okay and what isn't. If you're reflecting on AI feedback but making your own changes, you don't have to worry.

For all these reasons, I have focused on AI feedback since spring 2023 when I started serving as a volunteer advisor on the nonprofit app MyEssayFeedback. For four semesters, I've invited students to reflect on AI feedback, and they have overwhelmingly told me they find it useful. Still, I don't see AI feedback as a replacement for human feedback; I assign peer review and tutor visits too.

In my mind, AI feedback will never be enough because chatbots never experience our writing. They convert our draft into numbers and use a giant chatbot formula to spit out other numbers that convert to the feedback we read. That text might help us think about what we want to say, but we've still got to share the draft with humans and see how it lands. Getting a response from a person who's spent time reading or listening and gotten something out of what we wrote can be validating and energizing in a way chatbot feedback will never be.

Two very easy ways to get AI feedback

1. Sometimes the simplest approach to AI feedback works just fine. When I don't have the energy to think through exactly what kind of feedback I want, I will upload a draft to a chatbot and literally say "Give me feedback." The results are often useful and encouraging. If the feedback isn't quite what I'm looking for, my disappointment often helps me articulate what kind of feedback I do want, and I follow up with a more specific request. (We'll talk more about this below.)
2. The other easy way to get feedback is to copy and paste a feedback prompt into a chatbot along with your draft. I'm going to recommend the prompt I know best, the one I worked on intensively with my colleagues on the [Peer & AI Review + Reflection project \(PAIRR\)](#) in spring 2025. We drafted the prompt to incorporate core principles of writing pedagogy, and we tested it extensively. It tells the chatbot to describe two strengths and two possible areas for revision, quoting from your draft along the way.

The PAIRR prompt is designed to help you clarify what you want to communicate to readers. It's supposed to help you find meaning in what you're doing, rather than just telling you if your essay conforms to expectations or predicting what grade you'll get. We invite you to [read and comment on the full prompt](#).

To try it out, visit the [PAIRR Writing Feedback Demonstration bot](#), hosted on the nonprofit educational platform PlayLab.ai (note that PlayLab makes your chats visible to the PAIRR team).

Or you can copy and paste our feedback prompt into any chatbot. See the [Which AI tools should I consider?](#)

Follow these steps:

1. Copy and paste [the PAIRR prompt](#) into the chatbot (I know, it's long).
2. Copy and paste or upload your draft.
3. Optional: add your assignment instructions and/or rubric. (Our default prompt includes a basic common rubric for high-priority elements in writing-focused classes.)
4. Consider the feedback and respond.

Continuing the conversation to get more out of AI feedback

Even after years working with AI, my first impulse is still to take the first result from a chatbot as the best it can do. But a core principle of AI literacy is iteration: pushing the chatbot toward a better result.

Chatting back in response to AI feedback supports AI literacy and your development as a writer and thinker. Whatever you're thinking or feeling as you read the feedback, wherever you're stuck about what to do next, that might be something worth articulating, and the chatbot might give you something useful in response.

Seven strategies for getting more out of AI feedback:

- **Be frank and push back.** You don't have to sugarcoat it for AI or worry about how it will come across as you might with a human tutor, teacher, or peer. You can be more blunt. Disagree, object, or tell it when you're overwhelmed or frustrated—it's supposed to respond in a helpful and empathetic way.
- **Ask for clarification.** When anything isn't clear, ask the chatbot to explain. You can also ask for more depth, detail, or quotes from your essay to support its suggestions, or request examples of how to make the kind of revision it suggests.
- **Ask for a new version of the feedback.** If you're not satisfied, you can simply ask it to try again or give more insightful feedback. You can also request a different style of feedback based on your mood or preference, ask for the perspectives of a famous person you admire or a writer you've just been reading for class, ask how people with a particular life experience might respond, or ask for feedback that contradicts what it just said.
- **Ask about something it didn't address yet.** If the chatbot missed something in the assignment or rubric, something your instructor has asked you to work on, or something you're curious about, bring it up.
- **Get it to help you explore what you're not sure about.** Describe your uncertainties about what to do, and ask it to help you think through what's murky in your mind. You can also bring up something you're thinking of adding to the draft and ask for suggestions about how it could fit in.
- **Get it to suggest writing strategies.** Describe what you're struggling with and ask what writing process strategies you should consider.
- **Play with the bot.** Sometimes we freeze up because writing is hard or because we're nervous about sounding academic. Having a playful discussion of your draft might help you relax and explore and bring more of your authentic voice to your writing. Try asking it how you could work a random topic like whales into your draft. Or ask for feedback in the style of a pop culture character. If the chatbot answers silly questions seriously or gives you bizarrely contorted feedback, that's a good visceral reminder that these are prediction engines spinning text, not wise mentors—an AI literacy win.

Sample responses to AI feedback

Here are a few examples of the strategies described above.

Be frank and push back

- "That's seems condescending."
- "That seems biased against..."
- "Come on, didn't I already do that in this draft? I swear I did."
- "I'm not convinced I really need to make the thesis more specific."
- "I just don't think I can do that."
- "It just seems boring to add more examples as you suggest."

Ask for clarification

- "What do you mean by a more focused thesis?"
- "What's an example of splitting up a paragraph that's too long? How do you know where to split it?"
- "Can you share more quotes from the essay to show where I'm vague?"

Ask for a new version of the feedback

- "Try again."
- "Give me more insightful feedback."
- "Loosen up and explain it again."
- "It seems like you're dancing around something. Give it to me straight."
- "Be more supportive and help me see my strengths, but stick to the truth."
- "Can you explain it in a casual way like a friend would?"
- "What might James Baldwin / Gloria Anzaldua / George Orwell / Trevor Noah / Amy Tan say?"
- "What kind of feedback might I get from a reader who has experienced homelessness?"
- "Now give me a reasonable perspective that goes against what you just said."

Ask about something it didn't address yet

- "What about my introduction?"
- "I want to improve my sentence style."
- "My teacher suggested I work on integrating quotations."
- "How could my paper be more interesting?"

Get it to help you explore your uncertainties

- "Ask me questions about what I want to do with the conclusion."
- "I don't like the idea of cutting out the descriptions of all the kinds of solar energy although I agree that the reader might get lost in them. Can you help me explore why I don't want to cut them?"
- "That makes me realize I want to talk about the example of TikTok almost shutting down. Is there somewhere that could fit naturally?"

Get it to suggest writing strategies

- "I think my essay is all over the place, but outlines don't work for me. How else can I work on organizing it?"
- "I'm feeling overwhelmed. What are some strategies that could help?"
- "I can't think. How do I tell if I should take a break or keep pushing?"

Play with the bot

- "Okay, now tell me what possible connection my essay could have to popcorn / octopi / Batman / breakdancing. I want to put that in my conclusion."
- "What feedback might I get from Bart Simpson / Lisa Simpson / Spiderman / Ironman / Forest Gump / Grumpy Cat / Darth Vader / Aladdin's Genie" (Please insert your own favority character or comedian here—I'm not with it enough to guess what might appeal to you!)

Take a step back, reflect, and talk to a human

To get the most out of AI feedback, at some point we have to set it aside and use other strategies. Make sure you give yourself space to decide what you really think about AI feedback and what, if anything, you want to change in your writing as a result. Some teachers might ask you to write about your thoughts and plans after reading AI feedback; this is part of the Peer & AI Review + Reflection approach. If you're asked to do this, see [Writing a reflection about AI feedback](#) for possible topics.

Another way to get perspective on AI feedback is to seek out human readers. Sometimes I know just what I want to do with AI feedback, and I get right to it. Other times, I need to mull it over for a day. Still other times, the AI feedback makes me realize I need human input. If you're not sure if you agree with the AI feedback or you want more trustworthy advice, it could be a great moment to reach out to your teacher, a tutor, a classmate, friend, or family member.

Human feedback can also counteract the chatbot tendency to tell us what we want to hear. This is often called “sycophancy”; we could also call it flattery or “kissing up,” though chatbots aren’t hoping to get something from us. They’re just designed to produce things that humans tend to like. Chatbots can make us question our ideas if we ask them to, but more often, they tend to reinforce our inclinations and biases.

I ran into this problem recently when I asked Claude for feedback on a description of an upcoming workshop for teachers. I enjoyed the cycle of feedback and revision so much; Claude gave suggestions and kept telling me that each draft was a big improvement. I felt great when it celebrated the changes I was making because I too thought I was on the right track. The result was a workshop description that listed every point I was going to make and every activity I had planned. When I sent the description to the workshop organizer, she wasn’t impressed and gently suggested I cut out the overwhelming detail. At first, I was indignant and felt unappreciated. After a few hours I had cooled down and decided she was right. Claude had helped me talk myself into agreeing with myself, but the human workshop organizer was right that human teachers would feel overwhelmed on reading all the specifics of what I meant to do.

A next-level way to engage with AI feedback: Writing or adapting your own prompt

I’ve focused on the easiest ways to request feedback and then emphasized expressing your specific needs in the followup chat above. But once you’re familiar and comfortable with chatbot feedback, you’ll likely want to be more specific from the beginning. Writing your own initial feedback request or customizing someone else’s allows you to reflect on what kind of help works best for you in your writing process. It’s also a good way to practice prompting.

Here are some things you could include in a feedback prompt, though none of these are required:

- What’s important to you right now: what do you want to know or work on?
- What instructions has your teacher given you?
- What role do you want the chatbot to play? Examples might include supportive and understanding writing tutor or cheerleader, a rigorous critic, or a skeptical reader.

Here are a few sources of sample AI feedback prompts you can adapt or just browse for ideas:

- The [PAIRR prompt](#) mentioned above.
- The [MyEssayFeedback library](#) of feedback prompts teachers have written and shared (you can filter by whatever focus you want or sort by most popular)
- The PapyrusAI [prompt library](#) has prompts tested by faculty at the University of California, Irvine, each with a different writing process focus
- Professor [Mark Marino’s CoachTutor](#) prompt is playful and funny.
- The University of Sydney’s instructions for [Getting standards-based feedback](#) focus on meeting the assignment criteria.

If you do come up with a custom feedback prompt that works well for you, consider sharing it! Here’s a [site where students can post prompts](#).

Further reading

- [Getting the most out of AI \(prompting\): Describe what you want and keep trying](#)
- [Bad Idea: Now that We Have AI, We Can Ditch Human Feedback! Generative Idea: Combining Peer and AI Feedback](#) by Lisa Sperber
- [AI as Feedback Generator](#) by Ethan Mollick and Lilach Mollick
- [I, Cyborg: Using Co-Intelligence: How I used AI in my book about AI](#) by Ethan Mollick

This page titled [4.6: Getting the most out of AI feedback](#) is shared under a [CC BY-NC](#) license and was authored, remixed, and/or curated by [Anna Mills](#) (ASCCC Open Educational Resources Initiative).

4.7: AI for research assistance

If chatbots can be biased and inaccurate, then it's obviously a bad idea to use them to find information or sources, right? Well, these days, it's complicated.

You may have heard warnings against using ChatGPT like a search engine. You may have seen funny or not so funny examples of Google AI overviews giving wrong answers with footnotes right next to them. See [Don't trust AI: sometimes it makes things up](#).

However, I do find that in some circumstances they help me find an answer and check it much more quickly than a search engine. Especially now that most chatbots and other genAI systems can browse the Internet or databases and give links to real sources, there are ways they can be very useful as a complement to traditional search.

Table comparing traditional search to chatbot-assisted research

Traditional Internet or database search	GenAI-assisted research
You have to figure out the best search terms Pro: That means you have to think more about what you want. Cons: You might miss some alternate terms or it might take forever to try all possible combinations	You can just ask your question as you would naturally and the chatbot comes up with search terms. Pro: Less cognitive load and time Cons: You might not clarify what you're looking for in your own mind. Chatbot may misinterpret what you want and use the wrong search terms.
You go through a list of sources and decide which is relevant and credible.	GenAI searches and goes through the sources it finds and assesses which are credible and help to answer the question. It may choose sources that you wouldn't consider credible, like a Reddit post of someone's opinion or a site that advertises a product.
You read the sources and assess what the main points are and how they help you answer your question.	GenAI summarizes and compares the sources and answers your question very quickly. It may misrepresent what is in the sources.
Once you have searched and assessed what you found and arrived at an answer to your question, you'll have a good sense of what that answer is based on and will feel confident in its accuracy.	You'll get an answer quickly but won't know what it is based on. You'll have to research and read sources to see whether the answer is a good one and what supports it.
It might be prohibitively time-consuming to read and assess a large number of sources with different perspectives.	GenAI can analyze many sources quickly including multiple kinds of sources and perspectives on a topic.
Search may turn up lots of sources that are related but aren't precisely answering your question with all its details.	GenAI will provide an answer tailored to your question. For example, you can ask, "What are some little-known films I might enjoy if I love Steven Spielberg and dislike Martin Scorsese?"
	Quicker to give ideas for low-stakes questions where you just want to get your brain working and think of possibilities. Example:

This page titled [4.7: AI for research assistance](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills \(ASCCC Open Educational Resources Initiative\)](#).

CHAPTER OVERVIEW

5: Additional Resources

[5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI](#)

[5.2: Writing a reflection about AI feedback](#)

[5.3: Template Phrases for Critiquing AI Outputs](#)

This page titled [5: Additional Resources](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills](#) ([ASCCC Open Educational Resources Initiative](#)) .

5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI

Note: The sample essay below is shared with the student's permission under a [CC BY NC 4.0 license](https://creativecommons.org/licenses/by-nc/4.0/).

Lily Raabe

Professor Mills

College of Marin

ENGL 150

15 April 2024

How Dangerous Is AI?

Is there a safe way to use and develop AI, or will it inevitably be our downfall? The article “[You Can Have the Blue Pill or the Red Pill, and We’re Out of Blue Pills](#)” by Yuval Harari argues that AI, specifically large language models, could use its mastery of human language to destroy democracy, culture, and possibly humanity. Although the argument is compelling, it lacks evidence to support his claims and does not provide a strong rebuttal to the counterargument that AI potentially has limitless benefits.

In the article, Yuval Harari discusses the rapid development of AI and its potential negative impact on our future. The article begins with a shocking and somewhat unsettling statistic presented by Harari. He cites a 2022 study where 700 leading experts and researchers in artificial intelligence were asked about the possible future risks of a powerful AI currently in development. Half of the participants said they believe there is a 10% or greater risk of human extinction, and if not extinction, then a similarly permanent and severe disempowerment caused by AI. This is a startling fact- companies are putting significant amounts of money into something experts believe has a 10% chance of permanently hindering us or wiping us out completely. Harari argues that tech companies are neglecting to protect consumers from this potentially dangerous product and urges that we slow down the rate of development so society can prepare for the future of AI.

Harari follows this with another issue concerning him: AI’s proficiency with language. Harari claims that language is the basis of humanity and that without it, we would not have laws, money, art, or science. So, he asks, what does it mean for humans if artificial intelligence has a better grasp on language than we do? How does AI’s ability to generate language influence stories, images, laws, and policies? Harari also argues that AI can exploit human weakness, bias, and addiction and form intimate relationships with people. Humans could easily be taken advantage of if their weaknesses or personal feelings are being preyed upon. He worries that an advanced knowledge of language combined with the ability to manipulate could be detrimental to humanity. He writes, “In games like chess, no human can hope to beat a computer. What happens when the same thing occurs in art, politics or religion?”.

Harari believes that AI’s ability to master language will impact human culture. He points out that what started with AI writing essays for students could quickly turn into “political speeches, ideological manifestos, holy books for new cults.” This, Harari argues, could rewrite human culture. He continues by saying that humans are in a cultural prism, or cocoon, and how we see reality is affected by our cultural lens. He worries that if AI shapes culture, the way we interpret our reality will negatively change. He provides an example of using AI algorithms in social media to support this claim. Harari explains that companies use AI to create algorithms to keep people more engaged with the content that they view on social media. He argues that AI can create illusions that do not align with reality because the algorithm will curate content based on biases people have. The illusion he alludes to is that Donald Trump did not lose the 2020 election, and he believes people think this because of content fed to them by the algorithm. Harari claims that our current societal divide and the possible unraveling of democracy are due to social media and AI algorithms that help push propaganda.

After this, Harari provides a counterargument in his essay: “A.I. indeed has the potential to help us defeat cancer, discover lifesaving drugs and invent solutions for our climate and energy crises. There are innumerable other benefits we cannot begin to imagine. ” But he immediately rebuts it, saying, “But it doesn’t matter how high the skyscraper benefits A.I. assembles if the foundation collapses.” Harari argues that AI development is moving too quickly and that if we don’t create regulations and enforce the safe use of AI, we cannot safely take advantage of the benefits.

Harari follows this with a call to action. He demands control of AI development before there are extreme consequences and claims democracy is at stake if we continue down this path. He calls for world leaders to address the potential issues with AI and for

regulation around large language models. Finally, Harari suggests we prepare institutions for a world with AI and “to learn to master A.I. before it masters us.”

This was a compelling essay. Harari highlights potential issues that the rapid development of AI could cause. He emphasizes the importance of language to drive home the point that large language models hold more influence over us than we thought and that this influence could have far-reaching consequences for democracy and culture. He also brings attention to how urgently we need to place preventative measures so AI does not become uncontrollable. One weakness in the argument is Harari’s response to the counterargument. Harari acknowledges the counterargument but does not offer a convincing response. He does not explain why it is not worth the risk to find the benefits of AI or expand on his rebuttal. If AI has the potential to cure cancer or create lifesaving drugs, as he says, isn’t that worth exploring? If AI is being developed anyway, we may as well take advantage of it and use it to develop cures and solutions for worldwide problems. If we use AI in professional settings or in government funded programs to solve problems, it might also spur politicians to create regulations, which Harari wants.

The other weakness is that there is little evidence to support the claims. Harari cites a survey of experts at the beginning of the argument, but the rest is purely his opinion and does not provide evidence based on actual data. He offers no real justification for the idea that AI will use manipulative language and tactics that will lead to the fall of democracy; to convince us, he just repeats that idea with different phrasing.

Harari’s argument forces the reader to consider an uncomfortable prospect: AI could potentially eradicate humans. But is this really the case? AI has problems. It can be racist, show extreme bias, and make up sources or information, but could it lead to the fall of democracy, as Harari claims? We need more information about how exactly AI could use language to manipulate, exploit, or harm humans.

Harari uses borderline fearmongering tactics, writing, “However, simply by gaining mastery of language, A.I. would have all it needs to contain us in a Matrix-like world of illusions, without shooting anyone or implanting any chips in our brains. If any shooting is necessary, A.I. could make humans pull the trigger, just by telling us the right story.” But how would this be possible? Harari is correct that language is the basis of humanity and its culture, and I understand why a robot having mastery of language would be worrying. Still, his argument would be much more convincing if he showed how AI could wield language as a weapon against humanity. The claims seem exaggerated and unsubstantiated. Claiming that AI could manipulate humans into behaving in a way that kills us seems far-fetched. AI is not a free-thinking entity; it has no agenda. Why would it try to manipulate us? What would it gain from that? I wish he would provide evidence or explain his reasoning. I do not believe that the language Harari is using is purposefully manipulative but rather a tactic to force readers to recognize unregulated AI as a serious issue. However, his writing comes across as dramatic and hard to believe.

Despite this, Harari does a fantastic job of convincing readers to reflect on the possibilities of unregulated AI and the importance of language, culture, and human connection. His essay empowers readers to challenge corporations and world leaders and to think critically about the benefits they are promising from AI use.

Works Cited

Harari, Yuval, et al. “You Can Have the Blue Pill or the Red Pill, and We’re Out of Blue Pills.” *The New York Times*, 24 Mar. 2023, www.nytimes.com/2023/03/24/opinion/yuval-harari-ai-chatgpt.html.

Attributions

By Lily Raabe. Anna Mills changed the title (the original title was “Response Essay.”) Shared with permission under a [CC BY NC 4.0 license](https://creativecommons.org/licenses/by-nc-sa/4.0/).

This page titled [5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI](#) is shared under a [CC BY-NC 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/) license and was authored, remixed, and/or curated by [Anna Mills](#) ([ASCCC Open Educational Resources Initiative](#)).

- [Student Essay Critiquing a New York Times Article on the Dangers of AI](#) is licensed [CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/).

5.2: Writing a reflection about AI feedback

AI writing feedback can sound confident, authoritative, and academic and still be bad. Large language models like ChatGPT and the apps built on them are designed to predict likely next words, not to be right.

The downside of this is that we might be distracted or misled by AI. Jane Rosensweig, director of the Harvard Writing Center, explores this risk in her piece "[What Happens When a Novice Writer Asks ChatGPT for Editing Advice?](#)"

The upside is that being skeptical of AI feedback just might be empowering. We might clarify our own purpose in writing as we reflect on whether the feedback aligns with what we really want to say. Even if the feedback does resonate, it's still valuable to articulate why we like it. And more practice critiquing plausible-sounding AI suggestions never hurts in an era when AI will likely have an increasing presence in our personal and professional lives.

Below are some phrases that reflect possible ways we might respond to AI feedback. There will be countless others; these are offered as a beginning to give a sense of the possibilities. Please feel free to suggest additional phrases by commenting on this page.

Note: the quoted phrases followed by asterisks were generated by a large language model as detailed in the acknowledgments.

Reflect on what isn't clear to us about the AI feedback

- I'm not sure what the AI feedback could refer to when it says _____.
- The AI feedback advises that I _____, but I'm not clear whether that's necessary because the assignment says _____.
- I'm puzzled by the AI's suggestion to _____. How does that fit in a _____ essay where the instructor is asking for _____?
- I'm not sure what the AI advice to _____ means in practice. Does it mean _____?

Reflect on ways the feedback doesn't fit our meaning or purpose

- The AI feedback suggests _____, but I'm not sure this is what I want to do because _____.
- The AI feedback seems to assume that I am claiming _____ when in fact I meant to say _____.
- "The AI feedback suggests a misunderstanding of the section of my essay about _____. I was aiming to convey _____."*
- The AI feedback steers me toward _____, but that's not my priority because I think it's more important to _____.
- The idea of working on _____ in my essay seems reasonable, but it's more important to me to _____ because I want to make sure that readers understand _____.
- If I follow the AI suggestion to _____, I'm concerned that _____.
- "I'm not sure if the AI feedback fits the tone I was going for. It suggests _____, but I was aiming for a more _____ feel."*

Reflect on what we do like about the AI feedback and why it fits our purpose

- The AI feedback says "_____", and that helped me see that _____ needs work in my essay. I think the feedback is right about this because _____.
- I was already thinking _____ about my essay, so the AI feedback that _____ resonates.
- I've been concerned that _____, and the suggestion that _____ could help resolve that.
- The AI feedback suggests I clarify _____, so I am considering explaining that _____.
- "The AI highlighted _____, which I had overlooked. I think this could be important to include because _____."*
- I hadn't thought about _____ in the essay, but the feedback makes a good point that _____.

Reflect on new ideas we get as we read the AI feedback

- The AI feedback on _____ makes sense, but it made me realize I don't really want to keep the part where I say _____ after all. I really want to focus on _____.
- "Reading the AI feedback about _____ made me think of a new approach: I can talk about _____ in terms of _____."*
- "The AI comments on _____ reminded me of a related point I might want to add: _____."*

Next steps

Once we've decided what we think about the AI feedback, we can brainstorm what to do about it. We may want to write a revision plan or a list of writing tasks like brainstorming on a particular topic, rereading to clarify a point, researching to add evidence, or reorganizing a section. For more on revision strategies, see [the section on revision in my open textbook *How Arguments Work: A Guide to Writing and Analyzing Texts in College*](#).

AI feedback should complement, not replace, human feedback. After all, writing is a way to relate to others; it has meaning when we use it to communicate with humans. Writers often have breakthroughs when we talk through our drafts, our reactions to the AI feedback, and any uncertainties with our instructor, a tutor, or a peer.

Acknowledgments

The above template phrases followed by asterisks (*) were adapted from ChatGPT output responding to a "Template phrases for reflecting on AI feedback prompt," *ChatGPT*, 25 Sep. version, *OpenAI*, 8 Oct. 2023, <https://chat.openai.com/share/9c526e9b-b654-4c24-abb3-01a173026a63>. The remaining original phrases and the organizational structure are by Anna Mills and are shared under a [CC BY 4.0 license](#).

This page titled [5.2: Writing a reflection about AI feedback](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [Anna Mills \(ASCCC Open Educational Resources Initiative\)](#).

5.3: Template Phrases for Critiquing AI Outputs

Teaching critical AI literacy is now an essential aspect of digital literacy whether the purpose is to encourage critical perspectives on AI or to train students for future workplaces where they need to correct AI performance. The output of AI text generators like ChatGPT (also known as large language models, or LLMs) often sounds highly plausible and authoritative even when it is wrong, biased, or empty. The phrases below may help students practice identifying flaws in this polished text. Consider sharing a sample chat session as a web page and then inviting students to use these or similar phrases to annotate the outputs, perhaps using [Hypothesis](#) or [Perusall](#).

Note: the quoted phrases followed by asterisks were generated by a large language model as detailed in the acknowledgments.

Clarity

- This sounds plausible because _____, but it doesn't really make sense because _____.
- This sounds good, but it doesn't fit the purpose. What we are really looking for is _____.
- This doesn't clarify _____.
- "The AI is providing a surface-level answer without diving into the nuances of _____. "*"

Bias

- This is biased in favor of _____ because _____.
- This is biased against _____ since _____.
- This leaves out the perspective of _____.
- "This might be valid for _____, but not for the case we're discussing, where _____. "*"
- "It appears the AI has taken a literal approach, missing the cultural context of _____. "*"

Accuracy

- This is inaccurate because _____.
- This is outdated because _____.
- "The AI seems to have misinterpreted _____. "*"
- "While this might be true in some contexts, in the scenario of _____, it's misleading because _____. "*"
- The assertion that _____ has no basis in fact. _____ is a trustworthy source, and it says _____.
- _____ appears to be made up. When I searched for _____, I found _____.

Relevance and Focus

- "While the information seems relevant, it misses the point on _____. "*"
- This might be technically correct, but it doesn't consider _____. "*"
- "The AI appears to have focused on _____, but it's overlooking the importance of _____. "*"
- "This is a general statement that doesn't address the specifics of _____. "*"
- This seems to be a repetition of known facts rather than an insight into _____.
- "This is an over-simplification of the complex issue of _____. It overlooks _____. "*"
- "On the surface, this looks right, but it doesn't explain _____. "*"
- This is so vague that it is not useful in helping us think through _____.

Acknowledgments

The above template phrases followed by asterisks (*) were adapted from ChatGPT output responding to a "Template phrases for AI output critique prompt," *ChatGPT*, 25 Sep. version, *OpenAI*, 3 Oct. 2023, <https://chat.openai.com/c/a8b15d03-3...f-2014ced05511>. The remaining original phrases and the organizational structure are by Anna Mills and are shared under a [CC BY 4.0 license](#).

This page titled [5.3: Template Phrases for Critiquing AI Outputs](#) is shared under a [CC BY-NC 4.0 license](#) and was authored, remixed, and/or curated by [Anna Mills \(ASCCC Open Educational Resources Initiative\)](#).

AI use statement

Index

A

academic integrity

- 3: Academic Integrity
- 3.1: How do I know which use of AI is allowed and which isn't?

accusations of AI cheating

- 3: Academic Integrity

acknowledging AI

- 3: Academic Integrity
- 3.2: Acknowledging and Citing Generative AI in Academic Work

advice

- 5.2: Writing a reflection about AI feedback

AGI

- 1.4: AI copies patterns; it doesn't think

AI

- 1: Basics
- 1.1: Introduction
- 1.2: What is AI, and why are people talking about it so much?
- 1.4: AI copies patterns; it doesn't think
- 1.5: AI can generate decent-sounding text. Do we still need to learn to write?
- 2.1: Don't trust AI- sometimes it makes things up
- 2.2: Don't trust AI to cite its sources
- 2.3: Don't trust AI with sensitive information; check the privacy policy
- 2.4: Don't trust AI- it's biased
- 3.2: Acknowledging and Citing Generative AI in Academic Work
- 4.5: Which AI tools should I consider?
- 4.6: Getting the most out of AI feedback
- 5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI
- 5.2: Writing a reflection about AI feedback
- 5.3: Template Phrases for Critiquing AI Outputs

AI and academic integrity

- 3.1: How do I know which use of AI is allowed and which isn't?

AI and human extinction

- 5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI

AI and labor

- 2.5: Does using AI do harm? If so, should we stop using it?

AI and learning loss

- 3: Academic Integrity

AI and privacy

- 2.3: Don't trust AI with sensitive information; check the privacy policy

AI and writing

- 1: Basics
- 1.5: AI can generate decent-sounding text. Do we still need to learn to write?

AI apps

- 4: AI Strategies

AI assessment scale

- 3.1: How do I know which use of AI is allowed and which isn't?

AI assistance

- 4: AI Strategies

AI basics

- 1: Basics

AI bias

- 2.4: Don't trust AI- it's biased
- 5.3: Template Phrases for Critiquing AI Outputs

AI cautions

- 2: Cautions

AI cheating

- 3: Academic Integrity
- 3.3: How can I protect myself against false accusations of using AI?

AI consciousness

- 1.4: AI copies patterns; it doesn't think

AI detection

- 3: Academic Integrity
- 3.3: How can I protect myself against false accusations of using AI?

AI doom

- 2.5: Does using AI do harm? If so, should we stop using it?

AI energy use

- 2.5: Does using AI do harm? If so, should we stop using it?

AI error

- 2.1: Don't trust AI- sometimes it makes things up

AI ethics

- 2: Cautions
- 2.4: Don't trust AI- it's biased
- 2.5: Does using AI do harm? If so, should we stop using it?
- 5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI
- 5.3: Template Phrases for Critiquing AI Outputs

AI explanation

- 1.2: What is AI, and why are people talking about it so much?

AI feedback

- 4: AI Strategies
- 4.6: Getting the most out of AI feedback
- 5.2: Writing a reflection about AI feedback

AI for learning

- 4: AI Strategies

AI for students

- 4: AI Strategies

AI hallucination

- 2.1: Don't trust AI- sometimes it makes things up

AI harms

- 2: Cautions
- 2.4: Don't trust AI- it's biased
- 2.5: Does using AI do harm? If so, should we stop using it?

AI help

- 4.6: Getting the most out of AI feedback

AI in education

- 3: Academic Integrity
- 4: AI Strategies
- 5.2: Writing a reflection about AI feedback
- 5.3: Template Phrases for Critiquing AI Outputs

AI literacy

- 1: Basics
- 1.3: How do chatbots come up with text?
- 1.4: AI copies patterns; it doesn't think
- 2: Cautions
- 2.2: Don't trust AI to cite its sources
- 2.4: Don't trust AI- it's biased
- 4: AI Strategies
- 4.4: Getting the most out of AI (prompting)
- 4.5: Which AI tools should I consider?
- 4.6: Getting the most out of AI feedback
- 5.3: Template Phrases for Critiquing AI Outputs

AI mirage

- 1.4: AI copies patterns; it doesn't think
- 2.1: Don't trust AI- sometimes it makes things up

AI orientation

- 1: Basics

AI policies

- 3: Academic Integrity
- 3.1: How do I know which use of AI is allowed and which isn't?

AI policy

- 3: Academic Integrity
- 3.1: How do I know which use of AI is allowed and which isn't?

AI problems

- 2: Cautions

AI regulation

- 2.5: Does using AI do harm? If so, should we stop using it?
- 5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI

AI revision

- 4.6: Getting the most out of AI feedback

AI risk

- 5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI

AI sentience

- 1.4: AI copies patterns; it doesn't think

AI statement

- 3: Academic Integrity

AI systems

- 4.5: Which AI tools should I consider?

AI text generators

- 5.3: Template Phrases for Critiquing AI Outputs

AI tools

- 4: AI Strategies
- 4.5: Which AI tools should I consider?

AI tutor

- 4: AI Strategies
- 5.2: Writing a reflection about AI feedback

AI tutoring

- 4: AI Strategies

AI use statement

- 3: Academic Integrity

AI weirdness

- 1.4: AI copies patterns; it doesn't think

AI writing

- 4.6: Getting the most out of AI feedback

Anna Mills

- 2.2: Don't trust AI to cite its sources

Anna Mills and Maha Bali

- 2.2: Don't trust AI to cite its sources

APA

- 3.2: Acknowledging and Citing Generative AI in Academic Work

artificial intelligence

- 1: Basics
 - 1.1: Introduction
 - 1.2: What is AI, and why are people talking about it so much?
 - 1.3: How do chatbots come up with text?
 - 1.4: AI copies patterns; it doesn't think
- 2.2: Don't trust AI to cite its sources
- 3.2: Acknowledging and Citing Generative AI in Academic Work
- 4.5: Which AI tools should I consider?
- 4.6: Getting the most out of AI feedback
- 5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI
- 5.2: Writing a reflection about AI feedback
- 5.3: Template Phrases for Critiquing AI Outputs

B

bias in AI images

- 2.4: Don't trust AI- it's biased

botshit

- 2.1: Don't trust AI- sometimes it makes things up

C

chain of thought

- 4.4: Getting the most out of AI (prompting)

chain of thought prompting

- 4.4: Getting the most out of AI (prompting)

chatbot

- 4.6: Getting the most out of AI feedback

chatbot bias

- 2.4: Don't trust AI- it's biased

chatbots

- 1.2: What is AI, and why are people talking about it so much?
- 2.2: Don't trust AI to cite its sources
- 4.5: Which AI tools should I consider?

ChatGPT

- 1.2: What is AI, and why are people talking about it so much?
- 2.2: Don't trust AI to cite its sources
- 4.5: Which AI tools should I consider?

citation

- 2.2: Don't trust AI to cite its sources
- 3.2: Acknowledging and Citing Generative AI in Academic Work

cite

- 2.2: Don't trust AI to cite its sources

citing AI

- 2.2: Don't trust AI to cite its sources
- 3: Academic Integrity
- 3.2: Acknowledging and Citing Generative AI in Academic Work

Claude

- 4.5: Which AI tools should I consider?

copilot

- 4.5: Which AI tools should I consider?

copyright and AI

- 2.5: Does using AI do harm? If so, should we stop using it?

CoT

- 4.4: Getting the most out of AI (prompting)

critical AI literacy

- 1.3: How do chatbots come up with text?
- 1.4: AI copies patterns; it doesn't think
- 2: Cautions
- 2.2: Don't trust AI to cite its sources
- 2.4: Don't trust AI- it's biased
- 4.4: Getting the most out of AI (prompting)
- 5.3: Template Phrases for Critiquing AI Outputs

Critiquing AI

- 5.3: Template Phrases for Critiquing AI Outputs

D

data breach

- 2.3: Don't trust AI with sensitive information; check the privacy policy

data rights

- 2.3: Don't trust AI with sensitive information; check the privacy policy

datacenters

- 2.5: Does using AI do harm? If so, should we stop using it?

democratic oversight of AI

- 2.5: Does using AI do harm? If so, should we stop using it?

design justice

- 2.4: Don't trust AI- it's biased

detection

- 3.3: How can I protect myself against false accusations of using AI?

E

energy use

- 2.5: Does using AI do harm? If so, should we stop using it?

ethical AI

- 2.5: Does using AI do harm? If so, should we stop using it?

existential risk

- 2.5: Does using AI do harm? If so, should we stop using it?
- 5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI

F

false accusation of AI use

- 3.3: How can I protect myself against false accusations of using AI?

false positive

- 3.3: How can I protect myself against false accusations of using AI?

fast and slow thinking

- 1.5: AI can generate decent-sounding text. Do we still need to learn to write?

feedback

- 4.6: Getting the most out of AI feedback

future of AI

- 5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI

G

Gemini

- 4.5: Which AI tools should I consider?

genAI

- 3.2: Acknowledging and Citing Generative AI in Academic Work

generative AI

- 1.2: What is AI, and why are people talking about it so much?
- 1.4: AI copies patterns; it doesn't think
- 3.2: Acknowledging and Citing Generative AI in Academic Work

GPTZero

- 3.3: How can I protect myself against false accusations of using AI?

H

hallucination

- 1.4: AI copies patterns; it doesn't think
- 2.1: Don't trust AI- sometimes it makes things up
- 5.3: Template Phrases for Critiquing AI Outputs

honesty

- 3: Academic Integrity

How AI works

- 1: Basics
- 1.3: How do chatbots come up with text?

how chatbots work

- 1.3: How do chatbots come up with text?

I

intellectual property

- 2.5: Does using AI do harm? If so, should we stop using it?

intellectual property rights

- 2.5: Does using AI do harm? If so, should we stop using it?

introduction to AI

- 1: Basics
- 1.2: What is AI, and why are people talking about it so much?

IP

- 2.5: Does using AI do harm? If so, should we stop using it?

iteration

- 4.4: Getting the most out of AI (prompting)

iterative prompting

- 4.4: Getting the most out of AI (prompting)

L

large language models

- 1: Basics
- 1.4: AI copies patterns; it doesn't think
- 5.2: Writing a reflection about AI feedback
- 5.3: Template Phrases for Critiquing AI Outputs

learning with AI

- 4: AI Strategies

LLMs

- 1.4: AI copies patterns; it doesn't think
- 5.2: Writing a reflection about AI feedback
- 5.3: Template Phrases for Critiquing AI Outputs

M

Maha Bali

- 2.2: Don't trust AI to cite its sources

MLA

- 3.2: Acknowledging and Citing Generative AI in Academic Work

N

NotebookLM

- 4.5: Which AI tools should I consider?

P

p doom

- 5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI

p(doom)

- 2.5: Does using AI do harm? If so, should we stop using it?

Perplexity

- 4.5: Which AI tools should I consider?

playful prompting

4.4: Getting the most out of AI (prompting)

privacy

2.3: Don't trust AI with sensitive information; check the privacy policy

privacy policy

2.3: Don't trust AI with sensitive information; check the privacy policy

process tracking

3: Academic Integrity

3.3: How can I protect myself against false accusations of using AI?

prompt

4.4: Getting the most out of AI (prompting)

prompt engineering

4.4: Getting the most out of AI (prompting)

prompting

4.4: Getting the most out of AI (prompting)

prompting strategies

4.4: Getting the most out of AI (prompting)

R

reasoning models

4.5: Which AI tools should I consider?

revising with AI

4.6: Getting the most out of AI feedback

S

slow thinking

1.5: AI can generate decent-sounding text. Do we still need to learn to write?

T

text generation

1.4: AI copies patterns; it doesn't think

They Say / I Say

5.3: Template Phrases for Critiquing AI Outputs

training data

2.3: Don't trust AI with sensitive information; check the privacy policy

2.4: Don't trust AI- it's biased

Turnitin

3.3: How can I protect myself against false accusations of using AI?

U

understanding AI

1: Basics

1.2: What is AI, and why are people talking about it so much?

unreliable AI

2.1: Don't trust AI- sometimes it makes things up

V

version history

3.3: How can I protect myself against false accusations of using AI?

voice recognition

4.5: Which AI tools should I consider?

voice synthesis

4.5: Which AI tools should I consider?

voice typing

4.5: Which AI tools should I consider?

W

weird AI

1.4: AI copies patterns; it doesn't think

what is AI

1: Basics

1.2: What is AI, and why are people talking about it so much?

why write

1.5: AI can generate decent-sounding text. Do we still need to learn to write?

Writing

1: Basics

writing advice

5.2: Writing a reflection about AI feedback

writing and AI

1.5: AI can generate decent-sounding text. Do we still need to learn to write?

writing feedback

4.6: Getting the most out of AI feedback

5.2: Writing a reflection about AI feedback

writing instructions

1.5: AI can generate decent-sounding text. Do we still need to learn to write?

writing pedagogy

1.5: AI can generate decent-sounding text. Do we still need to learn to write?

writing to think

1.5: AI can generate decent-sounding text. Do we still need to learn to write?

Y

Yuval Harari

5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI

Yuval Noah Harari

5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI

Detailed Licensing

Overview

Title: [AI and College Writing: An Orientation](#)

Webpages: 40

Applicable Restrictions: Noncommercial

All licenses found:

- [CC BY-NC 4.0](#): 90% (36 pages)
- [Undeclared](#): 10% (4 pages)

By Page

- [AI and College Writing: An Orientation](#) - [CC BY-NC 4.0](#)
 - [Front Matter](#) - [CC BY-NC 4.0](#)
 - [TitlePage](#) - [CC BY-NC 4.0](#)
 - [InfoPage](#) - [CC BY-NC 4.0](#)
 - [Table of Contents](#) - [Undeclared](#)
 - [Licensing](#) - [CC BY-NC 4.0](#)
 - [1: Basics](#) - [CC BY-NC 4.0](#)
 - [1.1: Introduction](#) - [CC BY-NC 4.0](#)
 - [1.2: What is AI, and why are people talking about it so much?](#) - [CC BY-NC 4.0](#)
 - [1.3: How do chatbots come up with text?](#) - [CC BY-NC 4.0](#)
 - [1.4: AI copies patterns; it doesn't think](#) - [CC BY-NC 4.0](#)
 - [1.5: AI can generate decent-sounding text. Do we still need to learn to write?](#) - [CC BY-NC 4.0](#)
 - [2: Cautions](#) - [Undeclared](#)
 - [2.1: Don't trust AI- sometimes it makes things up](#) - [CC BY-NC 4.0](#)
 - [2.2: Don't trust AI to cite its sources](#) - [CC BY-NC 4.0](#)
 - [2.3: Don't trust AI with sensitive information; check the privacy policy](#) - [CC BY-NC 4.0](#)
 - [2.4: Don't trust AI- it's biased](#) - [CC BY-NC 4.0](#)
 - [2.5: Does using AI do harm? If so, should we stop using it?](#) - [CC BY-NC 4.0](#)
 - [3: Academic Integrity](#) - [CC BY-NC 4.0](#)
 - [3.1: How do I know which use of AI is allowed and which isn't?](#) - [CC BY-NC 4.0](#)
 - [3.2: Acknowledging and Citing Generative AI in Academic Work](#) - [CC BY-NC 4.0](#)
 - [3.3: How can I protect myself against false accusations of using AI?](#) - [CC BY-NC 4.0](#)
 - [4: AI Strategies](#) - [CC BY-NC 4.0](#)
 - [4.1: How do we tell uses of AI that help learning from ones that hurt?](#) - [CC BY-NC 4.0](#)
 - [4.2: AI for tutoring-style assistance](#) - [Undeclared](#)
 - [4.3: Reasons not to let AI brainstorm, write, revise, or edit \(in college\)](#) - [CC BY-NC 4.0](#)
 - [4.4: Getting the most out of AI \(prompting\)](#) - [CC BY-NC 4.0](#)
 - [4.5: Which AI tools should I consider?](#) - [CC BY-NC 4.0](#)
 - [4.6: Getting the most out of AI feedback](#) - [CC BY-NC 4.0](#)
 - [4.7: AI for research assistance](#) - [CC BY-NC 4.0](#)
 - [5: Additional Resources](#) - [CC BY-NC 4.0](#)
 - [5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI](#) - [CC BY-NC 4.0](#)
 - [5.2: Writing a reflection about AI feedback](#) - [CC BY-NC 4.0](#)
 - [5.3: Template Phrases for Critiquing AI Outputs](#) - [CC BY-NC 4.0](#)
 - [Back Matter](#) - [CC BY-NC 4.0](#)
 - [AI use statement](#) - [CC BY-NC 4.0](#)
 - [Index](#) - [CC BY-NC 4.0](#)
 - [Glossary](#) - [CC BY-NC 4.0](#)
 - [Detailed Licensing](#) - [CC BY-NC 4.0](#)
 - [Detailed Licensing](#) - [Undeclared](#)

Detailed Licensing

Overview

Title: [AI and College Writing: An Orientation](#)

Webpages: 40

Applicable Restrictions: Noncommercial

All licenses found:

- [CC BY-NC 4.0](#): 90% (36 pages)
- [Undeclared](#): 10% (4 pages)

By Page

- [AI and College Writing: An Orientation](#) - [CC BY-NC 4.0](#)
 - [Front Matter](#) - [CC BY-NC 4.0](#)
 - [TitlePage](#) - [CC BY-NC 4.0](#)
 - [InfoPage](#) - [CC BY-NC 4.0](#)
 - [Table of Contents](#) - [Undeclared](#)
 - [Licensing](#) - [CC BY-NC 4.0](#)
 - [1: Basics](#) - [CC BY-NC 4.0](#)
 - [1.1: Introduction](#) - [CC BY-NC 4.0](#)
 - [1.2: What is AI, and why are people talking about it so much?](#) - [CC BY-NC 4.0](#)
 - [1.3: How do chatbots come up with text?](#) - [CC BY-NC 4.0](#)
 - [1.4: AI copies patterns; it doesn't think](#) - [CC BY-NC 4.0](#)
 - [1.5: AI can generate decent-sounding text. Do we still need to learn to write?](#) - [CC BY-NC 4.0](#)
 - [2: Cautions](#) - [Undeclared](#)
 - [2.1: Don't trust AI- sometimes it makes things up](#) - [CC BY-NC 4.0](#)
 - [2.2: Don't trust AI to cite its sources](#) - [CC BY-NC 4.0](#)
 - [2.3: Don't trust AI with sensitive information; check the privacy policy](#) - [CC BY-NC 4.0](#)
 - [2.4: Don't trust AI- it's biased](#) - [CC BY-NC 4.0](#)
 - [2.5: Does using AI do harm? If so, should we stop using it?](#) - [CC BY-NC 4.0](#)
 - [3: Academic Integrity](#) - [CC BY-NC 4.0](#)
 - [3.1: How do I know which use of AI is allowed and which isn't?](#) - [CC BY-NC 4.0](#)
 - [3.2: Acknowledging and Citing Generative AI in Academic Work](#) - [CC BY-NC 4.0](#)
 - [3.3: How can I protect myself against false accusations of using AI?](#) - [CC BY-NC 4.0](#)
 - [4: AI Strategies](#) - [CC BY-NC 4.0](#)
 - [4.1: How do we tell uses of AI that help learning from ones that hurt?](#) - [CC BY-NC 4.0](#)
 - [4.2: AI for tutoring-style assistance](#) - [Undeclared](#)
 - [4.3: Reasons not to let AI brainstorm, write, revise, or edit \(in college\)](#) - [CC BY-NC 4.0](#)
 - [4.4: Getting the most out of AI \(prompting\)](#) - [CC BY-NC 4.0](#)
 - [4.5: Which AI tools should I consider?](#) - [CC BY-NC 4.0](#)
 - [4.6: Getting the most out of AI feedback](#) - [CC BY-NC 4.0](#)
 - [4.7: AI for research assistance](#) - [CC BY-NC 4.0](#)
 - [5: Additional Resources](#) - [CC BY-NC 4.0](#)
 - [5.1: Student Essay Critiquing a New York Times Article on the Dangers of AI](#) - [CC BY-NC 4.0](#)
 - [5.2: Writing a reflection about AI feedback](#) - [CC BY-NC 4.0](#)
 - [5.3: Template Phrases for Critiquing AI Outputs](#) - [CC BY-NC 4.0](#)
 - [Back Matter](#) - [CC BY-NC 4.0](#)
 - [AI use statement](#) - [CC BY-NC 4.0](#)
 - [Index](#) - [CC BY-NC 4.0](#)
 - [Glossary](#) - [CC BY-NC 4.0](#)
 - [Detailed Licensing](#) - [CC BY-NC 4.0](#)
 - [Detailed Licensing](#) - [Undeclared](#)