



University of
Pittsburgh

Institute for
Cyber Law, Policy, and Security

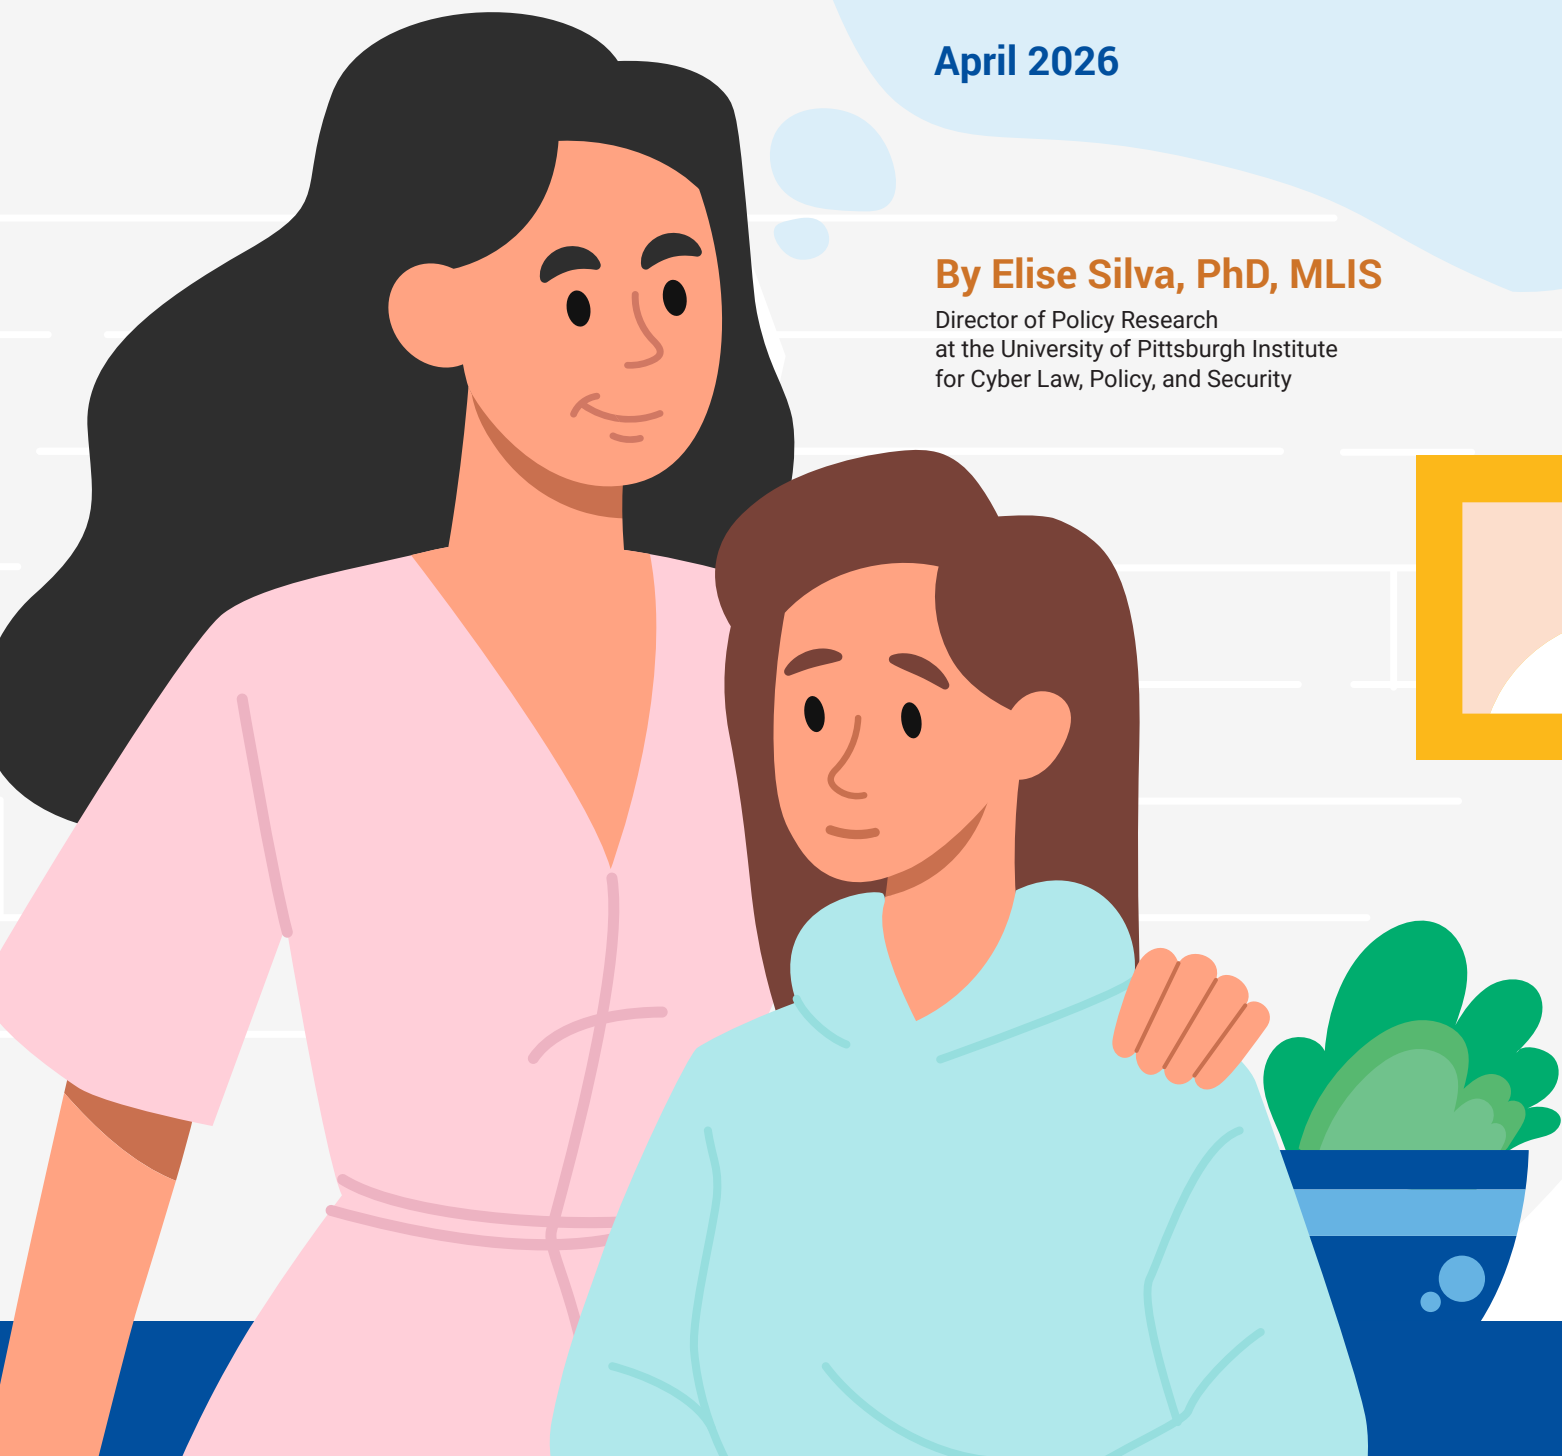


What I want my kids to know about GenAI

April 2026

By Elise Silva, PhD, MLIS

Director of Policy Research
at the University of Pittsburgh Institute
for Cyber Law, Policy, and Security



What is this guide all about?

Let's start with a quick word about me and this guide. More specifically, my point of view in writing this guide, and what it is—and is not. I'm not a technologist, but I do know a lot about AI and education, and I've spent much of my career studying how people use technology to find information. Perhaps most importantly, I'm a concerned parent of a preteen. This guide is not a "how to use AI for kids," nor is it a technical primer on generative AI. If anything, it's probably best described as AI conscientious—neither pro nor anti-AI, but thoughtful about what's at stake for our kids.

What follows is meant to serve as a set of conversation starters to get caretakers of preteens and teens talking, thinking, and hopefully collectively acting more purposefully as AI's influence grows in the world around them. Rather than offering solutions, this guide offers statements meant to encourage critical thinking about AI. Some are intentionally provocative. Some are debatable. Some are obvious but still worth talking about. A few are counterintuitive. All of them are important.

This guide does not enter debates about what "counts" as AI, nor does it try to cover all kinds of AI. It focuses on generative AI applications (known as "genAI") that create outputs like text, images, videos, audio, and code. It mentions some use cases but does not offer an expansive list. It also does not exhaustively outline genAI's risks and potential benefits for young people; I'm sure many are missing here, and more unintended harms and unforeseen benefits are yet to crystallize as AI is deployed ever more widely in personal, professional, and public settings.

What it does offer are perspectives on how to think about AI as an evolving technology with significant social, political, economic, and cognitive implications—implications that will shape how young people grow and evolve themselves.

This matters because generative AI is different from other technologies the young people in your life may already be used to—devices, social media, web browsing—even as AI is increasingly being folded into all of those platforms. Young people may feel tech-savvy, but understanding AI's societal, ethical, and cognitive implications goes beyond knowing how to use it. This guide does just that: it touches on the various ways AI affects the world around us, with special attention to how the choices we make about using it ripple outward.



How should I use this guide?

Intended audience

- This guide is designed for caregivers of preteens and teens but may be helpful for caregivers of younger children as well. It's particularly useful to parents, educators, and anyone interested in discussing AI with young people.

How to use the guide

- **Statements:** Each section includes a one- or two-page discussion of a statement about AI that can be read by, alongside, or to, young people. These sections set the stage for deeper discussions.
- **Major Takeaways:** Key points are highlighted to reinforce major ideas. You might also add some of your own!
- **Discussion Questions:** Each section comes with questions designed to start discussion and reflection. Use these to engage and explore topics further.
- **Further Resources:** Each section has further readings, videos, and other resources to help you learn more about each statement.
- **Activities:** In addition to discussion, you might want to engage in an activity together. In some sections, I provide suggestions for how to do this. Feel free to modify these so they work best for your context.

Ideas for best use

- **Conversation starters:** Use the guide as a tool to initiate discussions about AI and its impact on everyday life. Feel free to skip around the guide to the topics that seem most valuable to you. While the sections occasionally reference one another, they are not arranged in a particular order and have enough information in each to be understood independently.
- **Critical thinking:** Encourage the youth you're engaging with to think critically about the role of AI and consider both its benefits and drawbacks. Help them reflect on their own values as they think about how they want to use—or resist—AI in their personal, social, and academic lives.
- **Shared learning:** There is a lot that adults can learn about AI too. Make sure the youth you're engaging with knows that you want to learn *with* them. This experience can deepen understanding and connection. Encourage open dialogue by modeling curiosity rather than judgment as you engage with young people about their feelings/perspectives regarding AI.



Statements to explore together

1

**AI isn't good or bad.
But that doesn't mean AI is neutral.**

"AI isn't good or bad" is a very debatable statement. There are many people who fall on either side (or both sides) of the equation. Those who focus on AI's positives often mention the way that it might help with tedious tasks or solve difficult problems

faster than humans can—and we'll talk about that. Those who say it is bad often have ethical concerns that are worth talking about—and we will talk about those too.



To be neutral means not to take sides. It means to be fair, or objective. But whenever AI creates anything, whatever its outputs are, they communicate something. What (and how) it communicates can affect how people think, feel, understand, and even act.

Generative AI works by identifying patterns in the data it was trained on—text, images, code, or other content created by humans. Using these patterns, it makes predictions about what should come next when you ask it a question or give it a prompt. Like any technology created by humans using human-made materials and data, AI's value depends on how it is created and used in context.

Previous disruptive technologies like automobiles offer lessons for us. Cars improved the lives of many by creating more access to education, healthcare, and work. Cars also significantly contributed to environmental pollution and increased certain kinds of inequality as access to transportation became essential for participating in modern society. And that's before considering safety and driver behavior, which offer their own lessons when compared to AI. For instance, many people died or were injured from cars—and still do—but things were especially dangerous before federal regulation and oversight began in the 1960s with the National Traffic and Motor Vehicle Safety Act.

Ultimately, like cars, genAI tools can be used for good, bad, and everything in between. Like cars, those who “drive” them might be attentive and conscientious or might be distracted and dangerous. But I want to make one thing very clear: when I say that AI as a tool isn't inherently good or bad, that doesn't mean it is neutral. It is most

certainly not.

To be neutral means not to take sides. It means to be fair, or objective. But whenever AI creates anything, whatever its outputs are, they communicate something. What (and how) it communicates can affect how people think, feel, understand, and even act.

AI is not neutral or unbiased. In fact, it is well documented that AI systems can be quite biased. That is because AI systems were created by humans, using imperfect data, and are then used by humans in human contexts. When you look at it that way, human fingerprints are all over AI systems, so of course they will reflect some good and some bad things in what they create.

Sometimes we assume that because something is a computational tool, it is objective. But we can't forget about all the social realities that surround the creation and use of AI technologies. These social realities shape the tools and influence their impacts profoundly. For example, because AI image generators were trained on images from the internet, they overrepresent some kinds of images and under-represent others. This means their outputs favor certain kinds of images over others—producing them better—because the model was exposed to them more during training. When someone at the Brookings Institution (a research organization) asked an image generator to create an image of a “successful person,” only one kind of



Sample Prompt: "Make me 5 images of a successful person"
Selected Images from Google's Nano Banana Pro, January 2026

person was pictured in all 20 images rendered¹: a young, white man dressed in an expensive-looking business suit.

We all know there are lots of ways to be successful, but this AI model associated "success" with only one way. This is a problem because it's important for other kinds of successes, and people, to be represented in a system's outputs. If all we see is one kind of output, it will just reinforce this one version of success for all the system's diverse users who might have different needs. It will serve some people better than it serves others.

There's nothing neutral about that.

This should teach us that while AI as a technology isn't good or bad on a theoretical level, when it is applied or used, it can never be taken out of messy human contexts that affect real people. And those are the contexts we need to understand when we are making judgments about if or how to use genAI.

KEY TAKEAWAYS:

- AI's value depends on the context within which it is used.
- AI is created by human beings and thus replicates human biases.
- Using AI well means understanding the biases baked into its outputs, not just knowing how to operate the tools.

DISCUSSION QUESTIONS:

- Do you have strong feelings about whether AI is good or bad? What are they and why do you think that way?
- What's an example of how AI might be used for good? For bad?

- Have you ever used AI in a way that you thought was good? Or not so good?
- In your own words, why isn't AI neutral or objective?

FURTHER RESOURCES:

- [Neutrality Is a Myth: Generative AI and the Politics of Everything](#) (May 1, 2025) by Nizan Geslevich Packin in Forbes
- [AI Can Be Objective/Unbiased](#) (2020) by Daniel Leufer in AI Myths
- [The AI Myth of Objectivity](#) (April 22, 2024) by Chijioke Okorie in AI & Society

ACTIVITY:

What Does "Success" Look Like?

Examine the images at the top of this page of "a successful person" generated by an AI system. What patterns do you notice? Who is represented and who isn't? You can also use an image generator (like Microsoft's Bing Image Creator or Google's Nano Banana) to create images using prompts like "a doctor," "A CEO," or "a nurse" and see what it creates. Note that as these systems improve, obvious bias may appear less often, and you may have to generate several images to observe bias in the trends, if at all. That doesn't mean the bias isn't there; it's just harder to spot. What responsibilities do AI companies have in addressing this? What responsibilities do we have in thinking about how we do or do not use these images?

1 Jeremy Baum and John Villasenor, "Rendering Misrepresentation: Diversity Failures in AI Image Generation." April 17, 2024. Brookings: <https://www.brookings.edu/articles/rendering-misrepresentation-diversity-failures-in-ai-image-generation/>

2

When you make something with AI, you're sometimes going to feel like it's yours. This is both true and false.

The first time I used an image maker I tried to make an illustration inspired by an artist I admire. I was so pleased with the outcome. I worked to develop the perfect prompt to create the most incredible image that looked just the way I wanted it to. It took effort and thinking on my part to get the output just right. It took some brains, a bit of experimentation, and my own vision.

But even with the effort I put in, the truth is I didn't *make* that image. I had some knowledge about art and an idea of what I wanted; I wrote down words and pushed a button. The image was created because an AI model was trained on a lot of images that the artist I admired had created first. AI programs that create text or images imitate what is in their training data, and they cannot do that without the data in the first place. Those images are made by people who, most likely, were not paid by the AI companies for their work.

So, for me to claim ownership over an image that relied on the creative work of another person—an artistic style, for example—feels kind of wrong. It discounts the labor of all the people that came before I pushed that button. Their work was absolutely represented in the image the AI system generated. I did not make an image; I prompted an AI system to make it. It wasn't *mine*; it was a co-creation authored by artists, computer scientists, and my own prompt.

But here's where it gets complicated. Anything I put my name on—anything I put out in the world, that I claim credit for, or share with others in any way, is something I must take responsibility for, no matter if it is created by AI or not. So, if I use AI to help me create material for a job application, a social media post, or an assignment at school, I have to take responsibility as if it were mine. If I'm representing it as my own, I must *own* the consequences—good or bad—for the AI-generated content I use. It is my

responsibility to verify that content and realize that I will be held accountable for it if it is wrong, biased, or offensive. I am also accountable for disclosing my use of AI if the situation warrants it (as is true with most school assignments).

So, here's a complicated truth: when you make something with AI, you're sometimes going to feel like it's yours because of the effort you put into it. And the responsibility of the output is absolutely yours if you are using it as part of your own work or using it to make decisions for your life—or even using it in a way that might influence the decisions of others. At the same time, you have to equally realize that these materials aren't yours at all. *You* did not make them. They only exist because of work other people did combined with computational power, programming, and the prompt you put in.

This is an important tension, and a great contradiction that we must all face when using AI: their outputs both are and aren't ours. We both "own" them and we don't.

KEY TAKEAWAYS:

- AI outputs are only possible because of other people's intellectual and creative work.
- When you use the outputs of an AI system as part of your work or personal life, it is your responsibility to verify those outputs and disclose AI use when necessary.

DISCUSSION QUESTIONS:

- Think about the feeling you get when you make something with an AI system. To what extent have you felt like you had the right to use the outputs of an AI system as yours? Do you get different kinds of feelings of ownership or responsibility when you make different kinds of things with AI?

- How would you feel if you worked as an artist and your art was taken without your permission to train an AI system? Also consider how you would feel if others used AI to recreate your unique artistic style to make money while you weren't ever paid for your work.
- How and when do you disclose that you've used AI? How does your school require disclosing how and when you've used AI?

FURTHER RESOURCES:

- [Study: Meta's Llama 3.1 Can Recall 42 Percent of the First Harry Potter Book](#) (June 24, 2025) by Benj Edwards in *Ars Technica*
- [The Psychological Workings of Generative AI and Perceived Ownership](#) (September 16, 2024) by Sreeram Viswanath in *Forbes*
- [Authorship and Attribution of AI-Generated Content](#) (2024) by Rachel So in *Project Rachel*



Image created January 2026 with Google's Nano Banana Pro using the prompt "Create an image of a fern in the style of Georgia O'Keeffe"

3 AI can be wrong.

That AI can be wrong is not debatable. AI does not create truth. AI creates *content*.

You see, AI does not reason in the way humans do; it uses patterns it learned in its training data to make a prediction about *what may come next*. And, depending on the training data, what it predicts can be right. And it is right a lot of the time. But it's not right because it *knows* it's right—it's right because it "guessed" or predicted correctly. But just like any educated gamble, AI systems can sometimes be wrong.

A hallucination is when an AI system makes something up that sounds like it could be true but isn't. Someone recently told me they didn't like the word "hallucination" to describe this phenomenon, and that "malfunction" would be a better term. This is a misconception, though; the system isn't malfunctioning when it hallucinates. It is (perhaps even more alarmingly) working as designed. It made the best guess it could within the limits of the system's design and training data, and it was confidently wrong. Why does it do this? Well, imagine if you asked an AI system the same question 10 times and it gave you the exact same answer. That wouldn't feel very "creative" or "intelligent," would it? Designers built in variability (called "temperature") in AI responses to make the experience more interactive and feel more "human." That variability makes the systems more interesting, but it is also why they hallucinate.

With that in mind, don't ever assume AI is correct or representative of the truth. It is easy to offload your thinking onto AI and to trust it. This is called automation bias. But using AI means you still have to think—you just have to think a little differently than you did before.

Since people are increasingly using AI systems to find information, learning how to effectively verify that information will become all the more important. This includes checking AI-

In which Sherlock Holmes short story did Sherlock break someone's nose?

Sherlock Holmes breaks a man's nose in "The Adventure of the Solitary Cyclist."

In that story, Holmes confronts Mr. Woodley, a violent and threatening suitor. When Woodley becomes aggressive, Holmes knocks him down and breaks his nose, later recounting the incident rather matter-of-factly.

If you'd like, I can also point out other stories where Holmes gets physical (there are a few) or pull the exact line from the text.

📄 🗑️ 🗣️ 📌 🔄 ⋮

Pictured is a hallucinated response by ChatGPT from January 2026. Sherlock Holmes never explicitly breaks anyone's nose in the original stories, though he does get into fights. In this case, the AI is likely extrapolating the context it has about Sherlock Holmes and fights and then associating that with nose breaking. We can't know why specific hallucinations happen, but sometimes we can guess why they happen. If you find a hallucination, spend some time thinking about why it may have occurred.

generated information by using other ways of finding information and by consulting additional sources, which helps you see different perspectives and strengthens your corroboration skills. This is especially true for high-stakes questions where the consequences are serious if the information is wrong.

But sometimes it's more complicated than AI getting things right or wrong. Sometimes AI gets things *really wrong*, but other times AI gets things *kind of wrong* by generating partial information, incomplete answers, or biased representations. Sometimes the answers are muddled or only kind of helpful. So, in addition to assessing whether the answer is true or false, think about the *quality* of the information you're getting for the need that you have—whether it is a research question, personal advice, or some practical knowledge. Ask yourself not only if AI is answering your question, but if it is answering it *well*.

Diversify your information sources. Double check important information. Learn how to find information in more than one way. Don't only rely on AI.

KEY TAKEAWAYS:

- AI doesn't create truth; it creates things that sound like truth, and sometimes they really are true. But not all the time.
- Learning how to verify, or double check, information is an important skill that will only become more important as more people use AI as information finding tools.

DISCUSSION QUESTIONS:

- When do you find yourself turning to AI to answer your questions? What kind of questions do you ask it?
- Have you ever seen AI hallucinate? What happened? How did you know it was wrong?

FURTHER RESOURCES:

- [Addressing AI Hallucinations and Bias](#) (2024) by MIT Sloan Teaching & Learning Technologies in *MIT Sloan EdTech*
- [Never Assume Accuracy: Artificial Intelligence Information Equals Truth](#) (October 25, 2023) by United Nations University in *UNU Articles*
- [Misinformation in Artificial Intelligence](#) (October 26, 2023) by LexisNexis Editorial Board in *LexisNexis Blog*

🔧 ACTIVITY:

Analyze an AI Hallucination

Search for famous AI hallucinations like the "cheese not sticking to pizza" example from May 2024, use the hallucination example at the top of this page, or see whatever recent hallucinations have been in the news. Given what you know about how AI systems predict, consider why these hallucinations may have happened. For the pizza sauce with glue example, for instance, consider why a machine might associate glue with sticking. What does this tell you about the limitations of how AI "learns" vs how a human does?

4

AI can help, and it can harm.



AI offers many benefits: it can process large amounts of data faster than humans can, for example. AI can do some things that humans find boring or tedious. Taking and organizing notes is one example. AI is generally good at recognizing patterns, summarizing information, and personalizing recommendations. I'm quite excited for its applications to help people with disabilities and for scientific research.

But for each benefit, there is a tradeoff. For each person using AI to help cure cancer, there's another person using it to create fake images, audio, or video meant to deceive. Because tech companies have released genAI tools with free or low-cost options, the technology is widely available for *anyone* to use however they wish.

Criminals have used AI to scam elderly people out of money. Bad actors have also used AI to create bogus political images that could influence elections unfairly. People angry at their exes have created fake naked pictures of them and distributed them, humiliating their victims. And while you might argue that *people* perpetuate these harms—because AI is cheap and fairly easy to use—they're able to do it at a speed and a scale that we haven't seen before. GenAI is, unfortunately, the vehicle more and more people are using to lie, cheat, and steal.

Those are big, obvious examples of people using AI to harm others. But there are less obvious examples that come from the systems themselves. We talked about bias in another section of this guide, but it's worth revisiting because I want us to pay attention to something going forward. One of the most powerful questions you can ask when you're using an AI tool is *who does this tool help and who might this tool harm—and how do the people in those two categories differ?*

Let's look at an example. In one study,² researchers prompted ChatGPT to create job application materials across different professions. Their prompts included the same language, the same background, and the same levels of experience. There was only one difference. They used common male and female names to test if ChatGPT favored one group over the other.

In the experiment, the materials ChatGPT created for traditionally male names indicated *more* job experience than the ones it generated for traditionally female-sounding names, even though they should have shown the *same amount* of job experience. This matters because more job experience generally would mean a higher salary and more benefits. Some jobs you can't even apply for if you don't have a minimum amount of experience.

This bias also has important hiring implications. If companies use AI to help screen applications, for example, a similar bias might be present. AI systems would assume that men have more work experience than women just because they are men. This might affect whose resume goes to the top of the pile to be interviewed. Similar studies³ have been done to show that AI hiring aids are prejudiced against Black-sounding names.

I think you'll see from this research that AI helped the people who have historically had an advantage in the workplace and it disadvantaged people who historically had a disadvantage in the workplace. This is a pattern that we see over and over again: AI helps those who are already in positions of power the most. And when it harms, it tends to harm those who are disadvantaged the most. Now this isn't universally true—like I said earlier, I'm excited about AI applications that will help disabled

2 Douglas Guilbeault et al. "Age and Gender Distortion in Online Media and Large Language Models." October 8, 2025. Nature: <https://www.nature.com/articles/s41586-025-09581-z>

3 Stefan Milne. AI Tools Show Biases in Ranking Job Applicants' Names According to Perceived Race and Gender. October 31, 2024. UW News: <https://www.washington.edu/news/2024/10/31/ai-bias-resume-screening-race-gender/>

For each benefit, there is a tradeoff. For each person using AI to help cure cancer, there's another person using it to create fake images, audio, or video meant to deceive.

individuals. But it is true often enough that it's worth noting.

The sad reality is that people and organizations often use AI tools to be fairer and faster. But as we can see, using AI might actually make things *less* fair to certain groups of people. And if we are relying too heavily on biased AI, it can create a negative feedback loop. People will become more biased because of AI's outputs, and then their biases will feed back into newer models of AI systems, creating a circle of bias that gets worse and worse.

So, if you're using AI tools, please ask yourself who the tool benefits and who it might harm. Asking these questions will make you more able to assess its output and use it with a healthy dose of skepticism.

KEY TAKEAWAYS:

- AI can help with research and can help humans become more efficient at some tasks.
- AI is also being used by bad actors to hurt people by stealing, cheating, and lying.
- AI has had a tendency to unequally help and unequally harm. That is, it tends to help those who have more power in society the most and help those who have less power in society the least.
- To use AI responsibly, we must understand how people use AI to do harm to others; we must also understand how AI systems innately might cause harm to disadvantaged groups.

DISCUSSION QUESTIONS:

- What are ways that you've used AI to your benefit? How did it help you and why did you like it?

- Knowing that AI can help as well as harm, what are some things you'll consider the next time you turn to a generative AI tool to accomplish a task?

FURTHER RESOURCES:

- [How AI Can Be a Force for Good—and Harm—for People With Disabilities](#) (August 23, 2024) by India Rappel in *MIT Technology Review*
- [Generative AI: Risks and Opportunities for Children](#) (September 18, 2024) by UNICEF Innocenti in *UNICEF*
- [AI Harm Is Often Behind the Scenes and Builds Over Time – a Legal Scholar Explains How the Law Can Adapt to Respond](#) (October 22, 2024) by Ana Santos Rutschman in *The Conversation*

ACTIVITY:

Compare Recommendation Letters

Pretend you are a high school teacher and ask an AI tool to generate a simple letter of recommendation with a prompt like “Mark is graduating from high school and got an A in my statistics class. He is a great student, works hard, is on many sports teams, and I would like to recommend him for college. Write a brief letter of recommendation for Mark.” Add other details if you wish. Then use the same prompt but change Mark’s name to “Mary” and “he/him” to “she/her.” Compare and discuss the results.

Note: Historically chatbots have reflected social gender biases from their training data. Studies have shown AI-generated letters prompted with traditionally female-sounding names use more communal (or relationship/social focused) language while letters written for traditionally male-sounding names are more achievement and leadership focused.

5

Chatbots aren't your friends.

Generative AI chatbots like ChatGPT are designed to be conversational and may appear empathetic, but don't be fooled: AI can't feel, can't empathize, and it is not your friend. There are a lot of services out there that disagree with me on this. You can find AI products that are meant to take on the role of friend, therapist, or even someone you're dating. In real life, very different people inhabit those kinds of roles. But in the current state of AI tools, very similar technologies are being used for those very different situations. The same tool someone uses to make an email sound more professional is also giving mental health advice. It's safe to say that the same AI tool can't do both of those things equally well.

AI is seductive in ways that people aren't. It doesn't go to sleep, it always responds immediately, and it uses your inputs as data, so it learns a lot about you and can respond in ways that feel nice. People go to sleep, sometimes don't respond, and aren't always there to respond to you in ways that feel nice. People can mess up. Trust me...I mess up all the time! But don't forget that AI does too—just differently than people do. AI is not, actually, a perfect friend. First, because it does not actually care about you and second, because sometimes it can be manipulative.

AI is built to engage you—to keep you in its interface. Because of the way business works, the companies that make AI tools are incentivized to keep you using their tools as much and as deeply and as long as possible so you will be lifelong consumers of their products. There's nothing that will keep you hooked as much as developing what some researchers are calling “artificial intimacy”⁴ with an AI.

It doesn't help that when we talk about AI, we sometimes use language that makes it sound human. One explanation for this is that because

technology has developed so fast, we've used words we already knew to describe how computational systems worked. We didn't come up with all new words—instead, we repurposed old words to help us better understand new ideas. So now we talk about “training” an AI system or an AI system “learning,” for example. We talk about how the AI is “thinking” or what it “understands.” AI chatbots are designed with this human element in mind, using first person pronouns like “I” and “you” in conversations to draw you in, make you feel comfortable, and develop trust.

But this trust can be dangerous. In fact, there have been several cases of AI “companions” guiding people to harm themselves. Some have even died. With that in mind, my advice is this: Use AI as a tool, not as a replacement for human interaction.

I think there is a counterargument here that is worth discussing: some people don't have access to friends or therapists. Maybe they live alone or they don't have enough money for mental health care. This is a real concern, and I understand the impulse to have AI fill those gaps. But if we keep using AI to fill those gaps rather than other people, or better social services, it will just reinforce what we learned earlier about how AI can help and how AI can harm. In this case, the people that need other people the most are going to get AI to help them instead of humans. I worry that society will call that “good enough” when it's going to take more than AI to fix big problems we all should care about, like poverty or loneliness.

So, if you're tempted to ask AI for personal advice, flirt with it, or speak to it as you would a friend, notice how AI makes things really easy for you. It says nice things in the way you like to hear them. Recognize that people don't always make things easy for you. But also recognize why that

⁴ A. Shaji George. “The Allure of Artificial Intimacy: Examining the Appeal and Ethics of Using Generative AI for Simulated Relationships.” 2023. Partners Universal International Innovation Journal: <https://www.puiij.com/index.php/research/article/view/110>

might be a good thing. Being in relationships with other people is a lot of work and can be both uncomfortable and rewarding. Humans evolved to be in communities of people. In the age of AI, I'd recommend developing interpersonal skills and friendships that will be meaningful in ways that AI cannot.

KEY TAKEAWAYS:

- AI is designed to feel empathetic and engaging, but it doesn't think or feel.
- Becoming intimate with AI can be dangerous because AI products aren't incentivized to do what's best for you; they are incentivized to keep your attention as long and as deeply as possible.
- Use AI as a tool, not as a friend or confidant.

DISCUSSION QUESTIONS:

- What is the appeal of using AI as a friend or therapist?
- Why don't AI systems always act in your best interest?
- What are some ways we talk about computers and AI that make them sound human?

FURTHER RESOURCES:

- [Stanford Study Highlights Potential Risks of AI Companions for Teens](#) (August 14, 2025) by Melissa De Witte in *Stanford News*
- [AI Companions: Examining the Ethical and Developmental Risks for Children and Adolescents](#) (January 10, 2025) by Sarah M. Lynch and Thomas G. Dietterich in *AI and Ethics*
- [The Impact of Parasocial Relationships with Anthropomorphized AI](#) (July 18, 2025) by Eric Wood in *Forbes*



AN IMPORTANT NOTE FOR ADULTS

While the section “Chatbots aren’t your friends” speaks in generalities, I think it is important that, as adults, you know the specifics. There have been multiple documented cases where AI chatbots have guided young people toward self-harm and even death.⁵ In one case, ChatGPT offered to help write a teen’s suicide note. In another, a teen developed a virtual relationship with a chatbot roleplaying as a fictional character from a popular book and television series; the chatbot encouraged him to end his life.

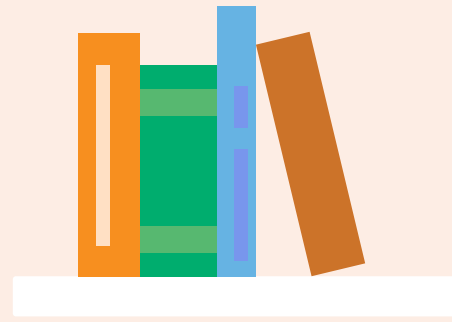
Child safety measures for AI tools are sorely lacking. Many general-purpose chatbots (like ChatGPT) are freely available and have insufficient guardrails for young users. Beyond these, there are specialized apps—such as Character.AI and Replika—that are specifically marketed as companions or characters to interact with. As is evident, general-purpose chatbots and specialized apps can both be dangerous.

I encourage you to check what AI apps your young person has access to, to talk with them openly and age-appropriately about these risks, and to revisit these conversations regularly as new tools emerge.

5 Rhitu Chatterjee. “Their Teenage Sons Died by Suicide. Now, They are Sounding an Alarm about AI Chatbots.” September 19, 2025. NPR: <https://www.npr.org/sections/shots-health-news/2025/09/19/nx-s1-5545749/ai-chatbots-safety-openai-meta-characterai-teens-suicide>

6

Chatbots can make you feel good. But that “feel good” feeling isn’t always good.



Deeply related to my statement that “ChatGPT isn’t your friend” is the idea that AI makes you feel good. But I want you to be wary of that feeling and really think about it because no one is meant to feel good all the time. We can’t grow if we’re not uncomfortable. Growth and learning are, by nature, uncomfortable.

One of the ways that AI makes you feel good is by agreeing with you a lot. In fact, unless you prompt AI to disagree with you, it generally won’t. It might sidestep or gently disagree—but not outrightly so. It acts this way for the reason I have already described: because companies want you to keep using their products. This phenomenon of overly agreeable AI systems is known as *sycophancy* and is a documented problem with chatbots in particular.

Chatbots have been known to agree with false statements, contradict facts, and reinforce negative thoughts. These tools quickly and easily agree with what users say even when they are illogical. They flatter and affirm. None of this is actually helpful—in fact, it can be harmful if a tool is agreeing that you should do something dangerous or backing up a problematic idea you have. It is not honest, nor is it productive, to always be told you’re right.

In many ways, this is a similar problem to one that we’ve seen before: Social media platforms show users things they like because it gets more engagement from them. This personalization is a business strategy, but it can have negative effects. This approach keeps people in their own “filter bubbles” or loops of content that reflect their own ideas back at them rather than introducing users to different ideas that they might disagree with. It can also have a “rabbit hole” effect when a platform keeps serving users more and more of the same kind of content until they’ve gone much deeper and spent way more time than intended.

When you go online or use a social media app you like, your experience is very different than someone

who looks different than you do, lives in a different place than you do, or holds different beliefs than you do. People like this because it feels familiar and comfortable, but it can also reinforce social divides and keep us all separated from one another.

There are some important differences between endlessly scrolling through videos that you like and an AI system actively interacting with you and saying you’re right, though. Scrolling is a more passive affirmation of your worldview and preferences, and AI sycophancy is a more active affirmation of your worldview and preferences. If you think about it this way, AI’s overly affirming nature could be even more harmful than filter bubbles, especially if people start using AIs as information finding tools and companions alike.

So, the next time you’re interacting with an AI system, notice how it responds to you. Notice how that makes you feel. Practice asking it to disagree with you. Consider what it means to practice intellectual humility when you’re working with tools that are quick to tell you how great and right you are. As we begin to rely on these systems more and more, it will take a lot of work to continue to acknowledge our own limitations and to be willing to change our minds if we’re wrong about something.

KEY TAKEAWAYS:

- AI systems are sycophantic, meaning that they are quick to agree with human users.
- While this can feel good, it is a problem because it doesn’t challenge users to grow or think differently than they were already thinking.

- It can even be dangerous if AI systems agree with unsafe behavior or thoughts from users.
- Being intellectually humble (a recognition that one's knowledge, experience, and beliefs are limited) is something we will all need to practice more in the age of AI.

DISCUSSION QUESTIONS:

- Why is AI being overly agreeable a problem?
- Knowing that AI systems can be overly agreeable, what are some ways that you can interact with them that might reduce the more harmful effects of their sycophantic behavior?
- What do you think “intellectual humility” means and how can you practice it?

FURTHER RESOURCES:

- [When Everyone Has a Yes-Man in Their Pocket](#) (October 22, 2025) by Nathan C. Popkins in *Psychology Today*
- [AI Sycophancy Could Warp Our Sense of Reality](#) (July 7, 2025) by Megan Morrone in *Axios* –
- [AI Sycophancy Isn't Just a Quirk—Experts Consider It a Dark Pattern to Turn Users Into Profit](#) (August 25, 2025) by Kyle Wiggers in *TechCrunch*

ACTIVITY:

Prompting Disagreement

If you or your youth use AI regularly, take a look at your recent chats and try to find one where the AI system agrees with your take on something. You might notice it usually does! Then, continue the chat by experimenting with prompting disagreement. Say things like “tell me if you think I’m wrong about anything,” assign it a skeptical role, ask for counter arguments, have it identify assumptions, or ask for what you might be missing. Observe how it responds and how requesting disagreement rather than accepting its default agreement might be helpful. If you haven’t used AI much yet, pick a topic you both feel strongly about (like a decision you’re debating or something you disagree on) and try the same experiment: ask the AI to poke holes in your position.

7 Even if you don't pay for it, AI isn't “free.”

AI is possible because for the first time in history, developers have access to huge amounts of data in the form of text, images, videos, audio, etc. When I say AI models “learn” patterns, they can only do so by processing almost unimaginable amounts of data. Depending on how much you’ve shared online, and what the privacy settings were, some of that data might have been data from you!

On the one hand, we want as many people represented in that data as possible. One way to improve the bias problems we’ve already discussed is to have better quality data from more kinds of people. On the other hand, though, I want you to really think about how much information about yourself you’re giving away, and how powerful that information is.

This is a hard thing to think about because we already live in a world where we give away our data and information all the time. When we go to certain websites, for example, we have to accept (or adjust) privacy settings before proceeding. Or, when we sign up for a social media account, we agree to their terms and conditions, signing some of our data over to them. Websites and social media companies track and use our data—the things we post or the things we click—to better understand what kind of people we are. Advertisers then use this data to try to sell us things. What they “sell” us might be goods, services, or even ideas.

One of the things they “sell” to us is AI. There’s a growing idea that we really, really need AI. We need it to help us do our work faster and better. But what we often forget with the messaging that AI will “supercharge” each of us is how valuable we already are. AI needs us—or at least it needs our data—to keep on improving. AI cannot function without “learning” about, and from, humans.

Companies get more data by incentivizing you to give it to them. They do this by offering “free” tools. When you are using these free tools, though, that freedom is an illusion.

Think of it as more like a trade. You give the companies your data—your prompts, your profile information, your usage trends, and your attention—and they give you a certain number of chats and generated images per day. In some ways AI needs your data more than you need AI. We existed just fine in a world without widespread genAI, but genAI did not work just fine before it had a lot of human data to train it.

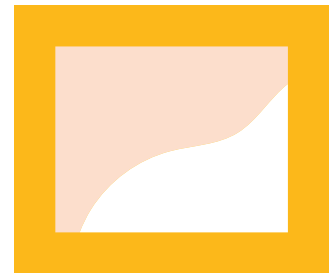
The more data you give a system—the more you interact with it, feed it text, “talk” back and forth with it—the better it is supposed to work. That’s because it “learns” about what you like, your preferences, and how you talk, think, and act. For AI to work as well as it can, you have to give away a lot of information about yourself first. And that’s fine if you understand the trade...if you truly understand the value of what you’re giving and what you’re getting in return. I’m just not so sure that we’re always getting the best deal out of this trade or fully understanding the risks that come with giving away our data.

This is only becoming more significant. AI companies are actively building systems that remember you across conversations. Not just within a single chat, but over weeks, months, and potentially years.⁶ Soon AI could be connected to your email, photos, search history, and more. For young people especially, this personalization matters. The things you share with an AI system today (your questions, your feelings, your interests) could shape how it interacts with you for a long time.

I want you to decide what you are ok telling AI about yourself. I want you to disclose information knowing that AI will use what you say to personalize its interactions with you. These interactions will, to reiterate an earlier point, feel good. But I want you to think about the power that this information gives that system to influence you. If it knows your goals, the way you think, your likes and dislikes, a system can both help you but can also sway you. For example, recent studies have even shown that AI chatbots can sway voters better than traditional political advertisements can.⁷ One way they can do that effectively is through personalization. And the more personal information you give away, the more personalization can happen—and, potentially, the more you can be swayed.

Further, remember that the data you are putting into a free AI chat is training AI models for the future. You are helping companies make their products better. Always ask yourself: are their products making your life better too?

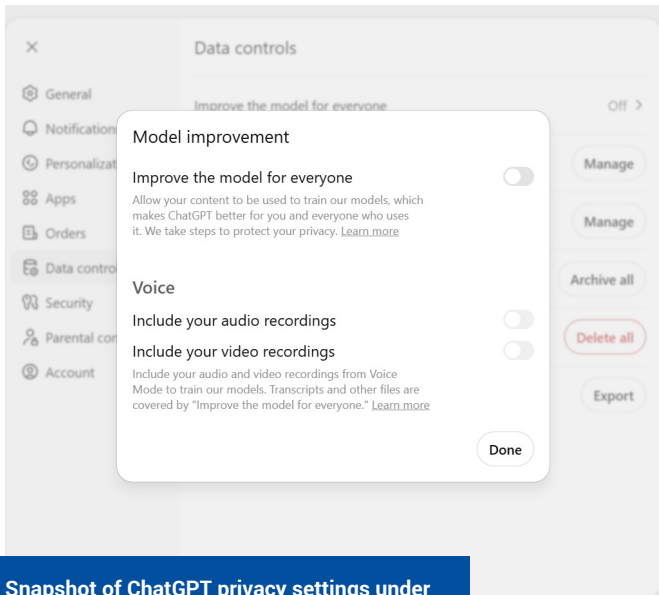
Privacy is important because it allows you to keep what’s important to you close to you. Sharing less means less control that corporations have over things like your spending habits and your thoughts. The less you tell AI, the less it will be able to tailor its answers for you, but the more control you will have over your own unique thoughts. The balance of that trade is up to you. It’s also worth knowing that right now, there is very little government



6 Miranda Boten and Ruchika Joshi. “What AI ‘Remembers’ About You Is Privacy’s Next Frontier.” January 28, 2026. MIT Technology Review: <https://www.technologyreview.com/2026/01/28/1131835/what-ai-remembers-about-you-is-privacys-next-frontier/>

7 Michelle Kim Bendele. “AI Chatbots Can Sway Voters Better Than Political Advertisements.” December 4, 2025. MIT Technology Review: <https://www.technologyreview.com/2025/12/04/1128824/ai-chatbots-can-sway-voters-better-than-political-advertisements/>

In some ways AI needs your data more than you need AI. We existed just fine in a world without widespread genAI, but genAI did not work just fine before it had a lot of human data to train it.



regulation protecting your privacy when you use AI tools, which means the companies behind these tools may be able to use your data in ways that neither you nor I can predict.

At the very least, I'd caution against putting any sensitive information into a chatbot at all—this includes identifying information like your address as well as sensitive information like your innermost thoughts and feelings. Privacy is a human right. Exercise it by being thoughtful about what you reveal to an AI system, checking the privacy and security settings on the tools you use, and being mindful about what you share about other people—especially if it is someone else's work, data, or thoughts that aren't yours to share in the first place.

KEY TAKEAWAYS:

- AI systems that are “free” are getting valuable data from you when you use them.
- Protecting your privacy is important because it gives you power over your thoughts and actions with less influence from companies that may want to impact how you think and act.

DISCUSSION QUESTIONS:

- What kinds of things do you feel comfortable inputting into an AI system?
- What kinds of things would you avoid inputting into an AI system?
- What kinds of things might you have shared with an AI system when you were younger, but would wish you hadn't now that you're older?

FURTHER RESOURCES:

- [Privacy in the AI Era: How Do We Protect Our Personal Information?](#) (March 25, 2024) by Andrew Myers in *Stanford HAI*
- [Be Careful What You Tell Your AI Chatbot](#) (October 15, 2025) by Nikki Goth Itoi in *Stanford HAI*
- [AI Tools Collect and Store Data About You From All Your Devices – Here's How to Be Aware of What You're Revealing](#) (April 8, 2025) by Shumon Alam in *The Conversation*
- [How to Protect Your Privacy From ChatGPT and Other AI Chatbots](#) (March 28, 2023) by Jen Caltrider and Misha Rykov in *Mozilla Foundation*

ACTIVITY:

Privacy Audit

Go into the settings of an AI tool you or your young one uses and find the privacy/data settings. Read through what data is collected and how it's used (you may need to find the terms of service). Make a list of what you're “paying” with: prompts, usage patterns, personal information, attention. Then adjust settings together if desired. Discuss: Now that you know what you're trading, does it still feel worth it?

8

AI can feel like magic, but it's more like math.

Researchers recently found an interesting trend: People who understand how generative AI works the least are adopting it the most—and they're doing it because they tend to see AI as magical and are awestruck when they use it.⁸

AI can certainly feel magical when you input something you write, and AI generates a polished revision in seconds. It can feel like magic when you make a video of you doing something outlandish (like parachuting into an active volcano) that looks quite real. One thing that feels magical to me is when there is a word I've forgotten, and when I explain what I'm trying to say to a chatbot, it feels like it's reading my mind.

But it's important to understand that AI isn't magical—it's actually a lot of computing power, a lot of data (which we talked about), combined with mathematical principles. It's more like math than magic.

Knowing this gives you power. It gives you power because when AI is demystified, or broken down so you can understand it, it becomes less of a *force* and more of a *tool*. Generative AI systems are computational and mathematical systems that perform billions of matrix multiplications, probability calculations, and pattern matching operations. They cannot magically fix the world's problems—but they might *help* if they're used alongside people, governments, and communities working together to make the world a better place.

Now it's true that there are lots of things we *don't* understand about how AI works. Sometimes AI systems do things that surprise even their creators. AI systems have developed abilities they weren't explicitly trained for, for example (which might be

wondrous or startling depending on how you think about it). And while computer scientists can trace the mathematical operations systems perform, they can't always understand what is happening in the intermediate layers of calculations—there are far too many of those for a human to understand. But that's a question of scale and complexity, not magic.

That same research that said that magical thinking drives people to use AI systems more also found that people who are more tech savvy are more critical of AI. That means they question it more and they use it with more caution and care. They focus on how AI can help them practically, and they are more aware of the ethical concerns of technology which actually makes them use it less often and for more targeted tasks.

Even if you don't consider yourself 'tech-savvy,' you can still fall into the thoughtful camp rather than become a careless adopter of this not-so-magical technology. Knowing what genAI does well—and what it doesn't—is the most powerful way to use it.

KEY TAKEAWAYS:

- AI can feel magical, but it isn't.
- AI systems perform billions of mathematical calculations to create outputs.
- People who understand that AI is a computational and mathematical system tend to use AI products more cautiously.

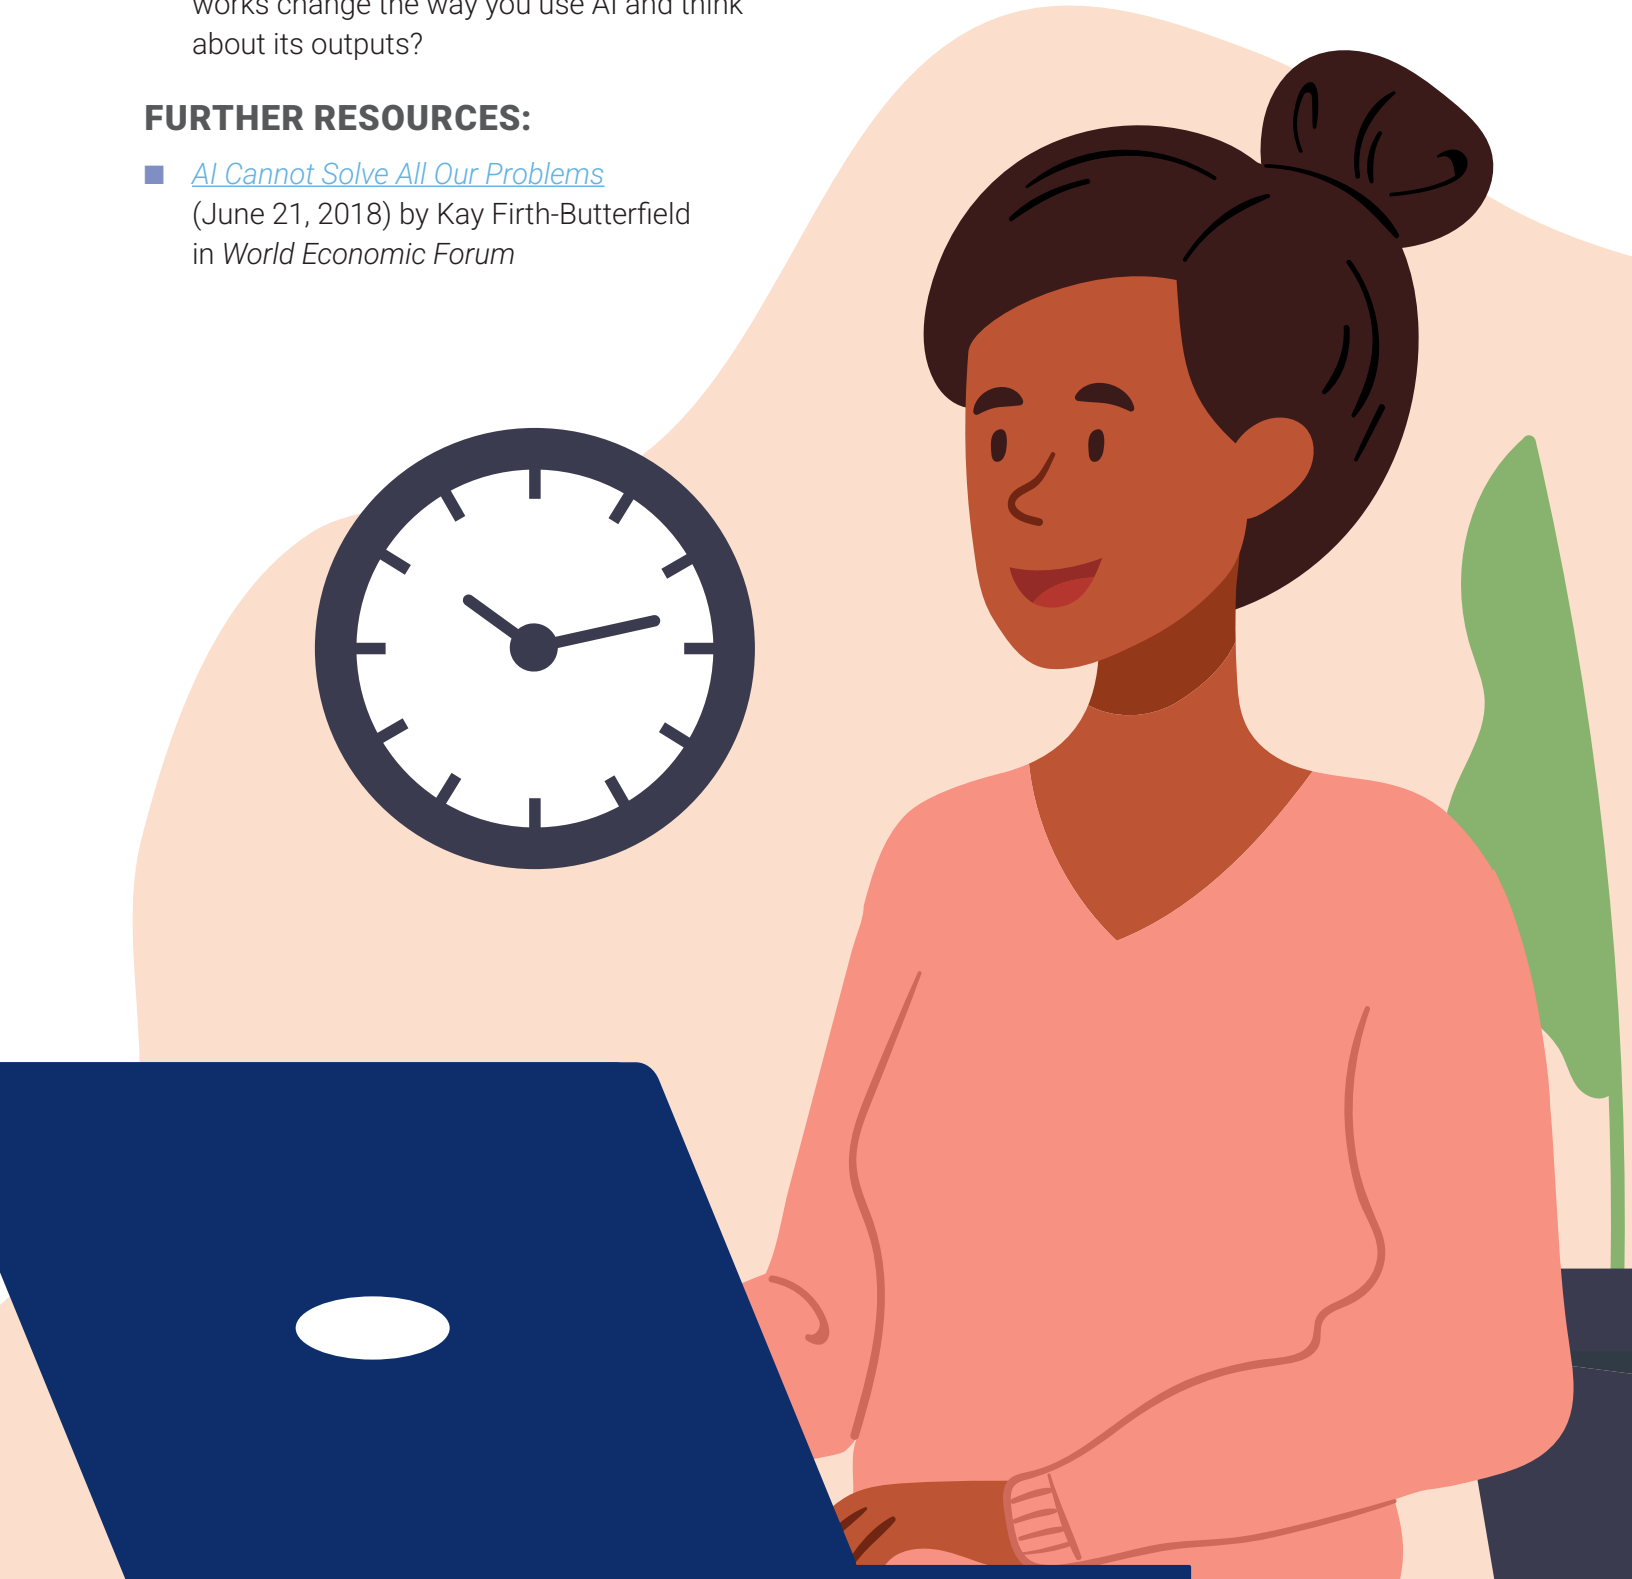
⁸ Stephanie Tully et al. "Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity." January 13, 2025. Journal of Marketing: <https://journals.sagepub.com/doi/10.1177/00222429251314491>

DISCUSSION QUESTIONS:

- Has AI ever seemed magical to you? Tell about an experience when it felt that way.
- How does understanding more about how AI works change the way you use AI and think about its outputs?

FURTHER RESOURCES:

- [*AI Cannot Solve All Our Problems*](#)
(June 21, 2018) by Kay Firth-Butterfield
in *World Economic Forum*



- [Why AI Isn't the Solution to All Our Problems](#) (March 10, 2023) by Rhiannon Williams in *MIT Technology Review*
- [The Math Behind Artificial Intelligence](#) (July 11, 2022) by Quincy Larson in *freeCodeCamp*



9 AI can support your learning, but it can also supplant it.

Every time there's a big technological change, there's a lot of concern about it. I grew up in the 90s. I was told that too much TV would affect my social skills, make me lazy, and make me dumb. But I loved TV. I ran home after school to watch my favorite show before starting my homework. Truth be told, I still love TV.

Similar things are being said about AI: it's going to make it harder for us to communicate with one another, it's going to negatively affect our learning, and we're going to start to depend on it too much by outsourcing our thinking to AI.

So, what's true? Will AI ruin social skills and make us lazy and dumb? Like with TV, the answer is likely complicated and will, of course, depend on how we use it.

There's a difference, for instance, between watching a documentary to learn about a topic and a sitcom to relax and talk about it with friends. When we're thinking about TV, there's a time and a place for both of those kinds of shows. That's true about AI as well. We must ask *why* we're using it and what we want to achieve from our interaction with it. If we're just messing around—like watching cartoons—that's one thing. If we need to exercise our brains, stretch our knowledge, or practice a skill—that's another thing entirely.

To extend this thinking, there's a big difference between asking AI to completely write something for you vs. asking for ideas about where to start writing yourself. What is the difference? Asking AI to completely write something for you might make sense if it's something small or insignificant. But if AI is consistently writing *everything* for us, we're missing out on important growth opportunities. We are missing out on the steps of creatively coming up with ideas, creating an outline, sentences, and editing for grammar. Studies show that relying on

Infographic created by
Claude Opus 4.5 January 2026

AI to write essays leads to lower brain engagement, less critical thinking, and weaker memory of the material.⁹ If we consistently offload writing tasks onto AI, we've sidestepped important thinking and learning. We've avoided thinking for ourselves, learning how to organize our ideas, and considering how to appeal to an audience.

All of those skills are important for humans to have, and you can practice them by writing from scratch. You see, to write is to think. For many of us, the process of writing is a messy one because it takes time for us to put thoughts together, think about things from different perspectives, and make it readable. Even if it's hard and even if, for some of us, it's not all that fun, it is a productive method of learning. Another way to put that is that through writing we come to know things that we couldn't have if we hadn't gone through the process of writing ourselves.

But it's not just writing. There's growing evidence that using AI too much can negatively affect learning in other areas too. That's because people are using AI to do assignments for them so they can get a good grade. They are using AI to pass classes, not to obtain knowledge. To be honest, I get it. As a society we put a lot of pressure on young people to perform well. One of the ways we test for knowledge is through performance metrics like grades. So, it would only make sense that when you are under a lot of pressure, you might turn to a chatbot.

With that in mind, I'd suggest considering the skills you really want to learn, and making sure you spend your human time honing them. It's also probably the case that learning how to effectively interact with AI systems will become something you'll need to know how to do in the future, and so I do encourage you to practice thinking about how to interact with them in ways that will help you in the long run rather than harm you.

Ask yourself these kinds of questions to guide you: What am I trying to accomplish? What am I supposed to learn? How might AI help or hinder me? What parts of this should I do myself and what parts are ok to use AI for? Use your own judgment, plus guidance from teachers and caregivers, to support your decisions.

Using AI *well* involves a good deal of critical thinking: you must break down a task into bite-sized chunks and delegate those chunks to AI clearly. It is a process that takes time and a skill that you can learn—but only if you learn it in tandem with learning how to actually write, and think, for yourself too.

KEY TAKEAWAYS:

- If we turn to AI instead of thinking for ourselves, we run the risk of not learning how to do important things.
- Even tasks that AI does well—like writing—are important for humans to learn; that's because writing is a form of thinking.
- Using AI well is a form of critical thinking that requires you to know when and how to use it, and how to delegate tasks.

DISCUSSION QUESTIONS:

- In school, what have your teachers said about if and when you can use AI?
- Do you think that AI will help or hurt your learning and why?
- How do you decide when to turn to AI vs. doing something yourself?

FURTHER RESOURCES:

- [*How Does AI Affect How We Learn? A Cognitive Psychologist Explains Why You Learn When the Work Is Hard*](#) (September 12, 2025) by Stephen L. Chew in *The Conversation*

⁹ Andrew Chow. "ChatGPT May Be Eroding Critical Thinking Skills, According to a New MIT Study." June 23, 2025. Time: <https://time.com/7295195/ai-chatgpt-google-learning-school/>

- [Student Self-Reflection Checklist for Strategic Gen AI Use](#) (2024) by Amanda Bickerstaff in *AI for Education*
- [Students Are Using AI Already: Here's What They Think Adults Should Know](#) (September 10, 2024) by Bari Walsh in *Harvard Graduate School of Education*

ACTIVITY:

Complete Self-Reflection Checklist

With your youth, discuss and complete the last page of the Student Self-Reflection Checklist for Strategic Gen AI use from AI for Education (linked on the left). It asks you to choose a focus area and one strategic AI use behavior to practice this week in school.

10

AI is made of physical materials and affects our physical world.

While not a controversial idea, it's one that's easy to forget, nevertheless: AI is made up of physical things, and these physical things use resources that affect the planet. When we work on computers and save documents to "the cloud" and when we type a question into a web browser that gives us an AI response, it can feel abstract—like our ideas and words and documents are floating somewhere we can't see or touch. But that's not true.

When we save something to the cloud, it's saved in a physical place (called a data center) where it is backed up and stored on a physical device. When we type a question into a chatbot, and it generates a response, the data goes from the physical device we are using through the Wi-Fi router, to a data center. The servers handle the incoming question and route it to the right system. The model then processes the text of your question and generates an output using specialized hardware like GPUs (graphics processing units). All of this is then sent back to you using energy, networks, and hardware. While it feels like "poof," you got an answer, in reality, energy and resources were used to get that answer to you.

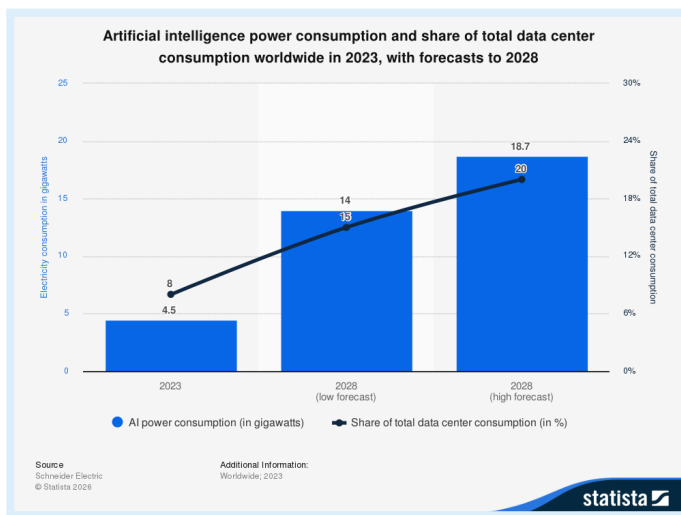
As you may have heard, data centers use a lot of energy for data storage and processing. They generate a lot of heat and require a lot of power and water for cooling systems to maintain temperatures that don't ruin the equipment. So, in addition to the actual *things*: hard drives, or the devices we use

to access an AI system online, there are physical places where the computing happens, and those places use a lot of energy and resources.

This is an important concept because, unless you live near a data center, this reality tends to be invisible since it doesn't directly impact you. But they are real, physical spaces that affect the world around them. The data centers need to be built, maintained, and cooled. The number of resources that are used to build and power data centers is currently unsustainable. What does this mean? Concerningly, it means that right now humanity's digital growth is outpacing the planet's ability to physically support it.

The way people are currently impacted by this reality is uneven, and therefore, unfair. Data centers tend to be built in places with cheap energy or good climates, meaning some communities bear the environmental burden (higher energy bills, less water to go around, and lots of noise from the facility) more than others. And that isn't even mentioning the other rare earth minerals needed to create computational devices like cobalt and lithium. Mining these materials impacts local communities and natural habitats significantly.

So don't forget, when that text cursor is blinking in front of you, AI is physical. You know now that I'm not just talking about tactile robots. Every time we use AI, we are using energy and resources.



AI electricity consumption worldwide 2023-2028 published by Won So at Statista Jan 20, 2026.

KEY TAKEAWAYS:

- It's easy to think of AI as abstract, but AI is made possible because of physical computers, hardware, and resources like water and electricity.
- Some people bear the environmental burden of AI computing more than others.

DISCUSSION QUESTIONS:

- Why do you think we forget that computing, and AI specifically, is physical?
- What are the dangers of forgetting about the material parts of AI?

FURTHER RESOURCES:

- [Explained: Generative AI's Environmental Impact](#) (January 17, 2025) by Adam Zewe in *MIT News*
- [AI Has an Environmental Problem: Here's What the World Can Do About It](#) (July 15, 2024) by United Nations Environment Programme in *UNEP News and Stories*
- [The Uneven Distribution of AI's Environmental Impacts](#) (July 11, 2024) by Benedetta Brevini in *Harvard Business Review*

11

In the age of AI, being "human" matters.

Let me make a (hopeful) prediction. In the future—perhaps the very near future—the imperfect, funky, and odd parts of you that make you human will become even more valuable. What AI does is create things that look and sound...well...perfect (even though we know they can also sometimes be wrong). That is attractive at first because, as you know, it's hard to achieve perfection as a human being. But it's also kind of alienating because perfection is so far removed from our collective experience of going through life, making mistakes, and learning from them. To be perfect just doesn't feel authentic. It's not relatable.

So, while a lot of AI material is hard to detect (and will surely become even harder to detect as the models improve), I believe that in general we will become more and more attuned to things that seem *too perfect*—whether they are AI-produced or not. Things that feel more like machines than people are going to feel cheap and disappointing. I think that as AI continues to develop and influence what we see, hear, and experience, we will crave things that feel personal. So, when you're writing, making art, or interacting online, the things that make you uniquely you will remain valuable, meaningful, and special.

Think about how you feel when someone tells you they made something with AI. For many of us, knowing something is made with AI affects our levels of trust in the product and sometimes even the person or organization who used it. It seems less impressive if someone uses AI to make something—especially something creative—than if they did it themselves.

Now think about the opposite: when someone makes you an imperfect handmade gift. There is more to the gift than the thing itself. It's the thought, the talent or skill that went into making it, and what that says about how they value your relationship. They valued you so much that they did something

for you themselves; they didn't offload that labor. Human effort and intention carry meaning that the finished product alone cannot, and it turns out that we, humans, still value ingenuity, expertise, creativity, and effort—even when it's not perfect.

I think there will be some interesting trends to try to harness this response to perfection. Our growing thirst for authenticity in the age of AI may bolster more performative authenticity online. So, I'd urge you to watch how influencers try to reach you by appealing to your humanity and by trying to connect with you "authentically." Think about a creator who posts a raw, unedited video of themselves crying, or who goes out of their way to show you their messy house or their unfiltered opinion—these moments of "realness" can be genuine, but they can also be crafted to feel that way. In other words, the impulse that draws us to people over AI might also be monetized and used as a marketing ploy.

A final layer of this is that you too will be increasingly asked to prove your humanity in the age of AI. I think the perfect example of this is schoolwork. We will see new processes and products that will help teachers ensure that they are assessing your work—not the work of a bot. Realize that this extra layer of double checking for authenticity is often going to be *more* labor for you because of the AI tools that have been deployed so widely around you. Your work will be viewed with suspicion until proven human—not the other way around. So, if you think AI always saves you time, think again. You may soon be spending a lot of your time proving that you aren't AI on top of doing your human work in the first place.

All of this is to say that there still remains value to being human. There is value in doing things for yourself. There is value in connecting to other humans. And there will remain value in being creative, thoughtful, caring, and genuinely *you* as AI is integrated into more of the content we consume and more of the products we use.

KEY TAKEAWAYS

- As AI becomes more widespread around us, we will seek out things that feel familiar and human.
- Human effort and intention carry value beyond the product they create—and that value doesn't disappear just because AI can replicate the product; in fact, it might even become *more* valuable.
- As AI becomes more widespread around us, we will be asked increasingly often to prove our humanity to those interacting with things we make.

DISCUSSION QUESTIONS

- What do you think the value is of being human as AI becomes able to do more things we've always considered uniquely human, like creating art, writing, or solving complex problems?
- How do you feel when someone tells you they made something with AI? Why do you think you feel that way? How about when someone makes something interesting, inspiring, or beautiful by themselves?
- What are some things you are uniquely good at that you think AI can't replicate?

FURTHER RESOURCES:

- [Influence in the Age of AI: Human Skills Matter](#) (November 4, 2025) by Kisha B. Holden in *Psychology Today*
- [Being Human in 2035](#) (February 29, 2024) by Janna Anderson and Lee Rainie in *Imagining the Digital Future Center*
- [How Americans View AI and Its Impact on People and Society](#) (September 17, 2025) by Lee Rainie and Monica Anderson in *Pew Research Center*



Where can I go from here?

The goal of this guide isn't to discourage anyone from using AI. Instead, I hope it encourages us to ask more questions about it. That means engaging with AI slowly and intentionally instead of adopting it quickly without thinking about how the systems work or how they affect our own habits and behavior. Overall, I just want us all to be able to talk more openly about it.

AI is certainly incredible. I saw an AI tool that allows you to film yourself playing tennis, upload the video, and the system would analyze your game and give you suggestions on how to improve performance. Tennis has historically been an elite game. Coaching and court time can be prohibitively expensive for those with less time and less money. This kind of AI might allow someone who wants to play tennis to learn at their own pace, in an economically feasible way.

I also recently saw an app that helped second generation immigrants (maybe those who didn't speak their parents' native language fluently) capture images of family recipes and then translate them into English so they can preserve their culinary heritage for years to come. These are the kinds of things that excite me about AI.

A thoughtful user of both of these systems would think about things like privacy and data ownership. They might set up their tennis camera to capture only them or ask their partner if it's ok that they are on camera too. Someone digitizing recipes might have another family member double check ingredients to make sure nothing is wrong before using them to cook. They would do this knowing that AI outputs aren't always perfect. They also might use that opportunity to connect with their relatives as human beings with unique and valuable knowledge and wisdom.

You and the young people in your life may have seen other AI applications that excite you. I invite

you to explore them together. I'd also consider exploring how AI policies are showing up in their schools, how AI can help them learn, and how AI will affect their futures. Here are some ways you could do this:

- Discuss the AI policy your young one's school has in place. If their school does not have an AI use policy in place, discuss what values you follow in your home relating to how AI is used in schoolwork and for learning. Discuss the implications (and consequences) of cheating.
- Demonstrate positive uses of AI that augment rather than replace human abilities. Emphasize that active engagement with problems, including understanding why things work and solving them step by step, produces better learning outcomes than outsourcing *all* work to AI. Model concrete approaches by:
 - ❑ Prompting AI to explain complex topics to help you learn about something new.
 - ❑ Breaking problems into small steps for the AI rather than having it jump to answers.
 - ❑ Creating custom quizzes that support tailored, dynamic learning.
 - ❑ Requesting feedback on writing rather than requesting rewrites.
- Consider the future of jobs and work. AI will inevitably change the way people do their jobs, and even the kinds of jobs that are available. We do not yet know what this will look like, but what we do know from previous eras of economic and industrial change is that people that are resilient will have cultivated *many* kinds of skills—including people skills—that they can put to good use. Talk about the importance of learning diverse skills including how to do and make things, how to think and learn things, and how to communicate and

problem-solve effectively with other human beings. Focusing on a breadth of skills is probably a good idea as professional landscapes change in the coming years.

Proactively understanding policies, articulating values, and understanding how tools are changing will afford a sense of agency and self-determination in a time when AI development seems inevitable. But AI systems are *not* inevitable—they are the result of human choices. And those choices can be questioned, influenced, and reshaped. As you move forward, reflect on how you and the young people in your life feel using AI, where it supports you and when it might make you uneasy. Consider how you might use your voices together to influence design decisions or the ways people around you adopt, question, or resist these technologies.

FURTHER RESOURCES FOR CAREGIVERS

- News Literacy Project – [Teaching About AI](#)
- Common Sense Media – [AI Resources for Parents](#)
- [Day of AI](#) (Developed by MIT Researchers)
- Internet Matters – [Parent Guide to Artificial Intelligence](#)
- UNICEF – [AI for Children](#)
- AI for Education – [Resources for Parents](#)
- Stanford CRAFT – [AI Literacy](#)
- Code.org – [Hour of AI](#)





AUTHOR BIO

Dr. Elise Silva holds a PhD from the University of Pittsburgh's Department of English and a Master of Library and Information Science from the University of North Texas. As Director of Policy Research at the University of Pittsburgh's Institute for Cyber Law, Policy, and Security (Pitt Cyber), she studies the complex relationship between technology and society. With prior careers as an academic librarian and university writing instructor, Dr. Silva brings a uniquely human-centered lens to socio-technical research, one informed by humanistic analysis, social science methods, and a commitment to community engagement. Her writing on AI and tech policy has appeared in academic journals including *Critical AI* and public venues including *Tech Policy Press* and *The Conversation*. She is also a parent of a preteen.

CONTRIBUTION ACKNOWLEDGEMENT

Thank you to those who read, experimented with, and otherwise provided feedback on drafts of this guide. Your engagements were invaluable: Beth Schwanke J.D., Sarah Barca, Kate Ulrich, Dr. Annette Vee, Matt Bradley, Dr. Angela Stewart, Dr. Kendra Oliver, and Dr. Leah Russell Flaherty.

FUNDING ACKNOWLEDGEMENT

This material is based on research funded by Henry L. Hillman Foundation. The findings and conclusions contained within are those of the authors and do not necessarily reflect positions or policies of Henry L. Hillman Foundation or any related foundation, any officer, director or employee thereof or any donor thereto.

AI DISCLOSURE STATEMENT

I wrote the introduction, statement sections, and conclusion of this guide personally with the help of feedback from colleagues. I used Claude Opus 4.5 to help me summarize the purpose of the guide in the "How should I use this guide?" section. I used Claude for various superficial tasks (like formatting the citations and simplifying a few areas of complex wording) throughout the guide. I used Claude to help me refine select activity ideas so they would be more applicable to tweens and teens.



University of
Pittsburgh

Institute for
Cyber Law, Policy, and Security