SUMMER 2026 ISSUE

Volume 1, Issue 2

Center for the Study of Race and Ethnicity in America

BROWN
Center for the Study of Race and Ethnicity in America

Natural Curiosities Pg. 10

Grow Your AI Literacy Pg. 28

The Trail Ahead Pg. 44



AN ARTIFICIAL INTELLIGENCE



FIELD GUIDE





TABLE OF CONTENTS

4 INTRODUCTION

6 GROUNDING

10 NATURAL CURIOSITIES

20 THOUGHTS FROM THE FIELD

28 GROW YOUR AI LITERACY

34 NAVIGATING THE FIELD TOGETHER

44 THE TRAIL AHEAD

48 APPENDIX

CENTER FOR THE STUDY OF RACE AND ETHNICITY IN AMERICA

A FAMILY FIELD GUIDE FOR NAVIGATING THE AI WILDERNESS



Artificial Intelligence (AI) is technology that learns, reasons, and performs advanced human-like thinking tasks. Its goal is to operate like a human brain. This means being able to brainstorm, categorize, predict, make decisions, and problem solve. We use it everyday, sometimes even without knowing it. AI technology is in our navigation systems, like GoogleMaps, our virtual assistants, like Siri and Alexa, and in streaming platforms like Netflix and Youtube.

Pause for a moment and consider:

IN WHAT WAYS IS AI SHOWING UP IN OUR DAY TO DAY LIVES?

AI is quietly changing how we think, learn, interact, and create — sometimes in really big ways! For kids who learn differently, AI tools can open up new ways to share ideas and show what they know. Families who speak multiple languages can use AI translation to feel more included at school or the doctor's office. But there's a catch: when AI does too much, like writing your book report, you might miss out on the learning that happens when you work through something hard on your own.

At Brown University's Center for the Study of Race and Ethnicity in America (CSREA), we're asking what kids and families think about these changes, and using what we learn to create AI-related resources for youth, caregivers, educators, and fellow researchers. In this field guide you'll find insights from youth and adults and activities for sparking conversations about AI. It is intended for caregivers and kids, ages 8 and up, to explore together. We hope you will learn together, talk together, think together, and become empowered AI users together.

01

GROUNDING



At CSREA, we believe knowing how to use AI is more than a technical skill, it's a deeply human one. As AI becomes a bigger part of life at home, at work, and in the classroom, knowing when and how to use it is becoming an essential skill for all of us. It takes agency, intention, and flexible thinking. We want you to be an active, informed AI user, not just a passive consumer. That means understanding when AI is actually helpful (and when it's not!) and knowing that your own judgement, cultural knowledge, critical thinking, and creativity should always take the lead. This is what sociologist and CSREA Director Dr. Prudence Carter calls culturally flexible AI literacy.

To kick off this work, we wanted to hear directly from people currently navigating AI's murky waters. We brought together middle and high schoolers, college students and professors, caregivers, educators, researchers, technologists, and community advocates for a series of community conversations — because the best way to understand something starts with listening.

CSREA: PLANTING SEEDS FOR EQUITY

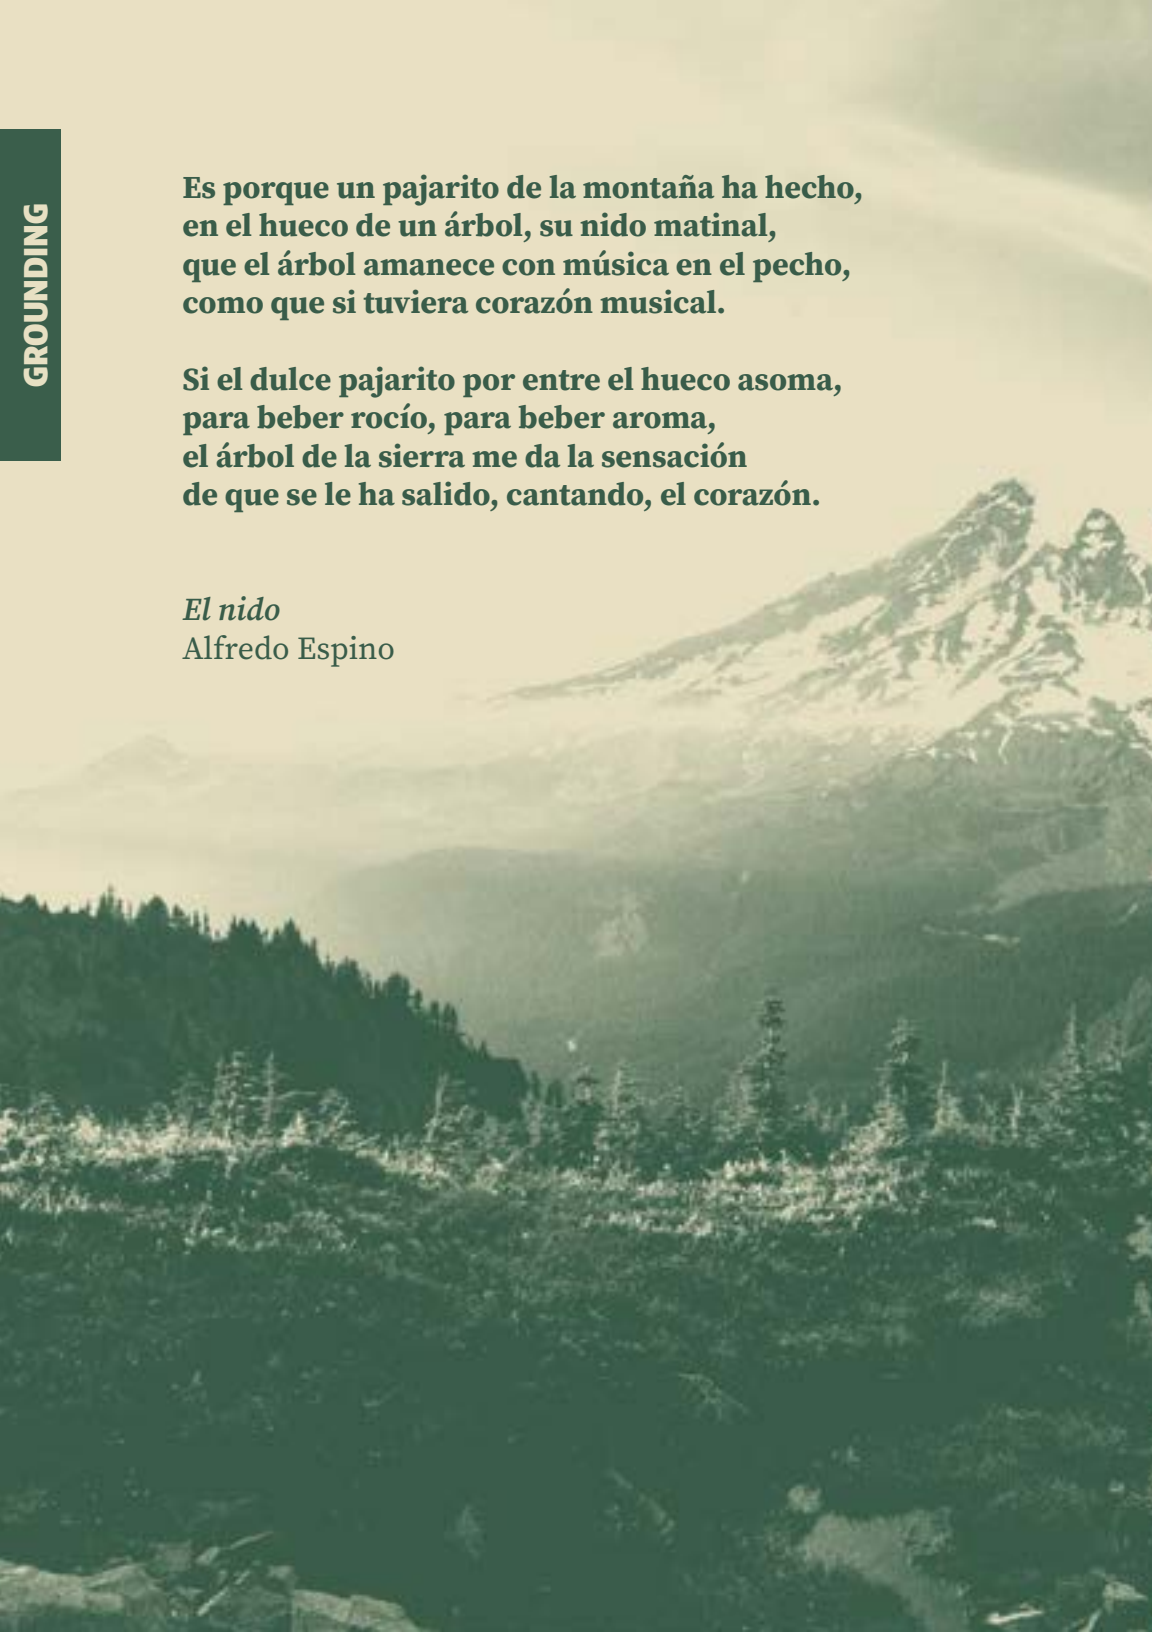
THE CENTER FOR THE STUDY OF RACE AND ETHNICITY IN AMERICA (CSREA) AT BROWN UNIVERSITY BRINGS PEOPLE TOGETHER – ACROSS CAMPUS, IN COMMUNITIES, AND AROUND THE WORLD – TO EXPLORE HOW RACE AND IDENTITY SHAPE OUR LIVES AND OUR WORLD.

THE CENTER'S WORK IS ALL ABOUT ASKING HARD QUESTIONS, HAVING HONEST CONVERSATIONS, AND BUILDING A FUTURE THAT'S MORE JUST AND INCLUSIVE FOR EVERYONE.

Es porque un pajarito de la montaña ha hecho,
en el hueco de un árbol, su nido matinal,
que el árbol amanece con música en el pecho,
como que si tuviera corazón musical.

Si el dulce pajarito por entre el hueco asoma,
para beber rocío, para beber aroma,
el árbol de la sierra me da la sensación
de que se le ha salido, cantando, el corazón.

El nido
Alfredo Espino



02

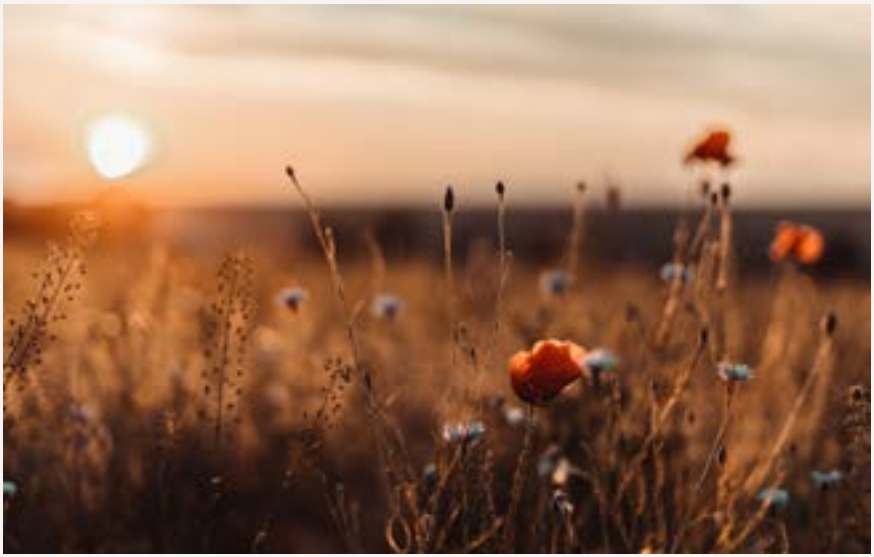
NATURAL CURIOSITIES



The youth we talked to are curious about AI. You may be curious, too. We've answered a few of their questions here.

WHO CREATED AI? HOW DID IT ALL START?

AI is rooted in machine invention. For centuries, engineers have experimented with making machines that could act autonomously, make predictions, and show reasoning — all to lighten human load. What we know as AI today, however, got its start in the 1950s with the invention of electronic computers. Scientists began to wonder if they could teach computers to think like humans. It started slow with simple tasks, but as time went on and computers became more advanced, they were able to perform more advanced tasks. The next big moment for AI was in the 2010s with the development of social media, online research platforms, and other big data sites. This sparked AI's rapid learning, and its ability to perform all the things we see it do today.



HOW DOES AI WORK? HOW DOES IT KNOW EVERYTHING?

All AI systems require three things. The first is vast amounts of data. Data is how AI learns, it's like its teaching or training material. It provides AI with millions of little examples, bits of complex information, and tons of images to analyze and learn from. The second is algorithms. Algorithms set AI's rules for learning; they're the instructions AI needs to follow in order to accomplish what we ask of it. It's people who decide or set AI's algorithms. The last is processing power, that's the computer's speed and ability to process all the data AI pulls from.

With these three things, AI can sort through huge amounts of data faster and more efficiently than humans can. It can recognize patterns that may be hard for us to find, make decisions by weighing all sorts of options, and identify and connect ideas using an almost endless knowledge bank. One of the most common types of AI is called a Large Language Model, or LLM. LLMs are the technology behind tools like Claude and ChatGPT. They're specially trained to understand and generate human language, meaning they can read your questions and write back in ways that sound a lot like a person. They've learned to do this by analyzing enormous amounts of text from books, websites, and other written sources.

Some AI systems take this a step further using something called Retrieval-Augmented Generation, or RAG. Instead of relying only on what it learned during training, a RAG system can go and look something up in real time. It's kind of like the difference between answering a question from memory versus being able to check a book while you answer. This helps AI give more accurate, up-to-date responses. Latimer is an AI system using RAG. It takes what LLMs come up with, cross references it with its information and resources, then adds to it to create an even richer response. RAGs also open up an exciting possibility: what an AI system retrieves depends on where it looks. That means the sources and knowledge bases we build for AI really matter.



CAN AI HELP US SOLVE BIG PROBLEMS?

It depends on how we use it! AI's abilities to sort through lots of data, automate repetitive processes, and analyze and predict patterns can help make problem-solving more efficient for the human beings who use it. For example, AI damage assessment tools help speed up the search process so first responders can focus more on rescues. Similar assessment tools are supporting medical professionals, too. AI can read scans to find patterns that help with diagnostics, giving doctors more time for patient care. In schools, some teachers use AI to help analyze student progress, identify learning gaps, and personalize standardized curriculum to meet student needs. When built and used thoughtfully, AI can be a helpful teammate, giving the people already tackling big challenges more time and tools to make a difference.

I HEARD AI USES A TON OF ENERGY AND WATER. IS IT ACTUALLY BAD FOR THE ENVIRONMENT?

Real talk: Yes! And it's worth knowing why. AI tech is housed in massive buildings called data centers. Data centers are kind of like a giant version of your laptop. The tech that powers AI needs huge amounts of electricity to run, and, just like your phone gets hot, it overheats too. So it uses enormous amounts of water to cool down. Data centers also pump pollution into the air, a lot like a factory. And because AI is always running, it's always having an impact.

Here's the part that really matters: AI companies often choose to build these data centers in low-income Black and Latine communities that are already dealing with air pollution, water pollution, and fewer green spaces. That's not an accident; it's a choice. The good news? Choices can be challenged. Learning about where data centers are built and who's most affected is a powerful first step. Communities have pushed back before, and yours can too.



Mt. Rainier National Park in Washington, United States

CAN AI COMPANIES JUST... USE MY PERSONAL INFORMATION WITHOUT ASKING?

Kind of — and it happens more than most people realize. AI learns from everything we do online: searching, posting, clicking, even just scrolling. It also has access to the cameras, microphones, and fingerprint scanners on our devices. All of that adds up to a lot of personal information; and once it's collected, most of us have no idea what happens to it next. Privacy agreements are often written to be intentionally confusing, which isn't an accident either.

That said, you're not powerless here. Asking questions like "What data does this app collect?" or "Do I actually need to sign up for this?" are small but real ways to protect yourself. Talking to an adult about privacy settings is a great place to start. Pushing for clearer, fairer privacy rules is also something young people and families can absolutely be a part of.



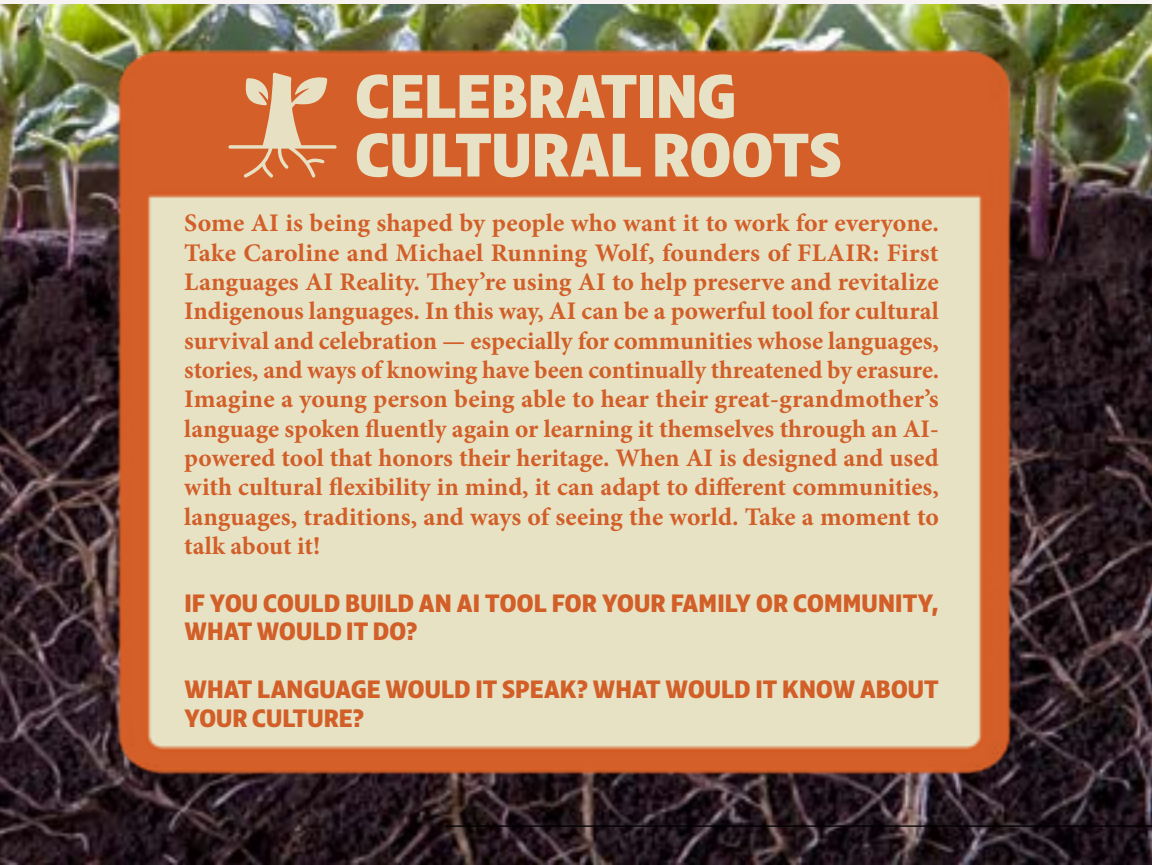


A data center in Biggleswade, England

WHAT WILL HAPPEN WITH AI IN THE FUTURE?

Honestly? AI is evolving fast and constantly changing. This makes it hard to predict where it is going, which is both exciting and a little nerve-wracking! What started as a tool used mostly by researchers in labs has become something most people interact with every single day, from homework helpers to music generators and the recommendations that pop up on your favorite apps.

While we don't really know what will happen next, we do know the people building AI are being very intentional about making it more powerful, harder to avoid, and more profitable. Sometimes they're moving so fast the rules or protections that keep us safe can't keep up. That means we need to be intentional about how we're using AI, too. "Intentional" is just a fancy way of saying we need to use it with thought and care!



CELEBRATING CULTURAL ROOTS

Some AI is being shaped by people who want it to work for everyone. Take Caroline and Michael Running Wolf, founders of FLAIR: First Languages AI Reality. They're using AI to help preserve and revitalize Indigenous languages. In this way, AI can be a powerful tool for cultural survival and celebration — especially for communities whose languages, stories, and ways of knowing have been continually threatened by erasure. Imagine a young person being able to hear their great-grandmother's language spoken fluently again or learning it themselves through an AI-powered tool that honors their heritage. When AI is designed and used with cultural flexibility in mind, it can adapt to different communities, languages, traditions, and ways of seeing the world. Take a moment to talk about it!

IF YOU COULD BUILD AN AI TOOL FOR YOUR FAMILY OR COMMUNITY, WHAT WOULD IT DO?

WHAT LANGUAGE WOULD IT SPEAK? WHAT WOULD IT KNOW ABOUT YOUR CULTURE?

SOMETIMES WHEN I CREATE AN AVATAR, IT LOOKS NOTHING LIKE ME. IS AI BIASED?

We all want our gaming persona to reflect aspects of our identity — from our hair, skintone, and gender, to our hobbies. The thing is, AI is not neutral. Not even close. Whether creating an avatar or generating an image, AI doesn't necessarily reflect everyone equally. AI learns by studying massive amounts of data, mostly from the internet. But here's the catch: that data was created by humans, and humans carry social biases — conscious and unconscious ideas about race, gender, sexuality, disability and more. When AI learns from biased data it doesn't just absorb and replicate the bias, it amplifies it.

Here's what that looks like in real life: Facial recognition software (AI that identifies people by their faces) has been shown to misidentify Black and brown faces at much higher rates than white men's faces. AI image generators, when asked to picture a "doctor" or a "scientist," often show white male figures. And when asked to portray "beauty," they often default to euro-centric beauty standards. LGBTQ+ people and people with disabilities are frequently either invisible in AI systems or reduced to harmful stereotypes.

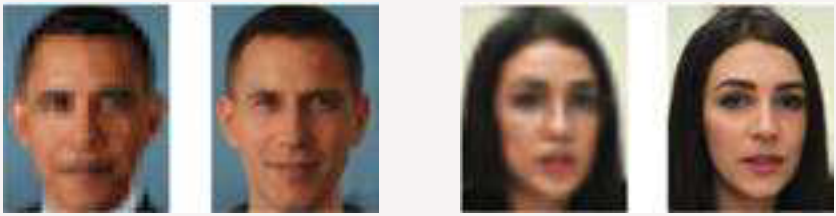


Image Resolution in Black and White: Behind the Pulse Controversy About Bias in Machine Learning
https://www.deeplearning.ai/the-batch/image-resolution-in-black-and-white

None of this is inevitable. This is how LLMs are designed and trained. And fortunately they can be challenged and changed. Researchers, activists, and young people are already pushing AI companies to draw on more diverse, representative data, test their tools for bias before they are released, and ensure that Black, Indigenous, Latine, LGBTQ+, and disabled communities are at the table when decisions about AI are made. When more diverse voices help build AI, the technology gets better and more just for everyone. You don't have to be an engineer to make a difference. Keep asking, "Who does AI see, and who does it leave out?"

HOW CAN I KNOW WHAT'S REAL? HOW CAN I KNOW WHAT'S FAKE?

As AI gets more advanced, this question gets harder to answer. But asking is the most important step! In general, learning more about AI and understanding its strengths and weaknesses can help you decide when to trust the information it gives you. Use your own intuition, and be skeptical like you would with any other source. Just because it has access to lots of information doesn't mean it has the ability to judge that information. If something does not sit right, dig deeper into it. Ask questions, and find community with others who are remaining critical and curious.



AI LEAVES A TRAIL

There are concrete things you can look out for as you strengthen your AI literacy. Start by paying attention to content you know is AI generated. There are certain patterns you'll see in text such as:

- **OVERUSE OF SENTENCE STRUCTURES LIKE, "... NOT X, BUT Z."**
- **IN LONGER TEXT, PARAGRAPHS ARE UNIFORM IN LENGTH AND REPETITIVE IN STRUCTURE**
- **SIMPLE, NEUTRAL CONCLUSIONS THAT OFTEN START WITH PHRASES LIKE, "IN SUMMARY..."**
- **EMOJIS AT THE END OF ALMOST EVERY SENTENCE**
- **FREQUENT USE OF THE EM-DASH (—) AND BULLET POINTS**

And in images and videos:

- **SMALL TEXT AND LETTERING ARE BLURRY OR DO NOT MAKE SENSE**
- **BACKGROUNDS WITH ODD AND UNCLEAR FEATURES**
- **FINE DETAILS, LIKE THE NUMBER OF FINGERS OR BLINKING PATTERNS, ARE UNREALISTIC OR CHANGE THROUGHOUT THE VIDEO**
- **OVERLY SMOOTH TEXTURES**

While these patterns do not always mean content is AI, frequent use or overuse can be your first sign to dig deeper. Observing details closely will help you start to notice the trails AI leaves in the wild.

IT'S SCARY THAT AI IS SO *HUMAN*. WHAT CAN I DO TO NOT LET IT MESS WITH MY HEAD?

It may seem like it sometimes, but AI doesn't have original thoughts, feelings, or ideas, and it doesn't have real social-emotional or interpersonal skills. That means AI can't replace real-life human communication and connection. It can be easy to lean on a resource that feels non-judgemental, but remember, AI isn't a friend, it is a tool! Ground yourself in what real friendship feels like and nurture your human relationships. In other words, stay connected to your people. Take digital breaks (touch grass!) and when you are online, be intentional with your choices.



03
THOUGHTS FROM THE FIELD



Kids and caregivers share similarities in their perspectives about AI use. They both talked about AI's popularity, its prevalence and overuse, and that it can negatively impact learning:

"People rely on AI. If you keep relying on AI and not learning something for yourself, and it comes down to not using AI, you're not going to be able to do it."

- 7th Grader

"It's everywhere. Now they've added it to Tiktok, you don't even get to agree to it. It's just there. You kinda have to just adapt to it."

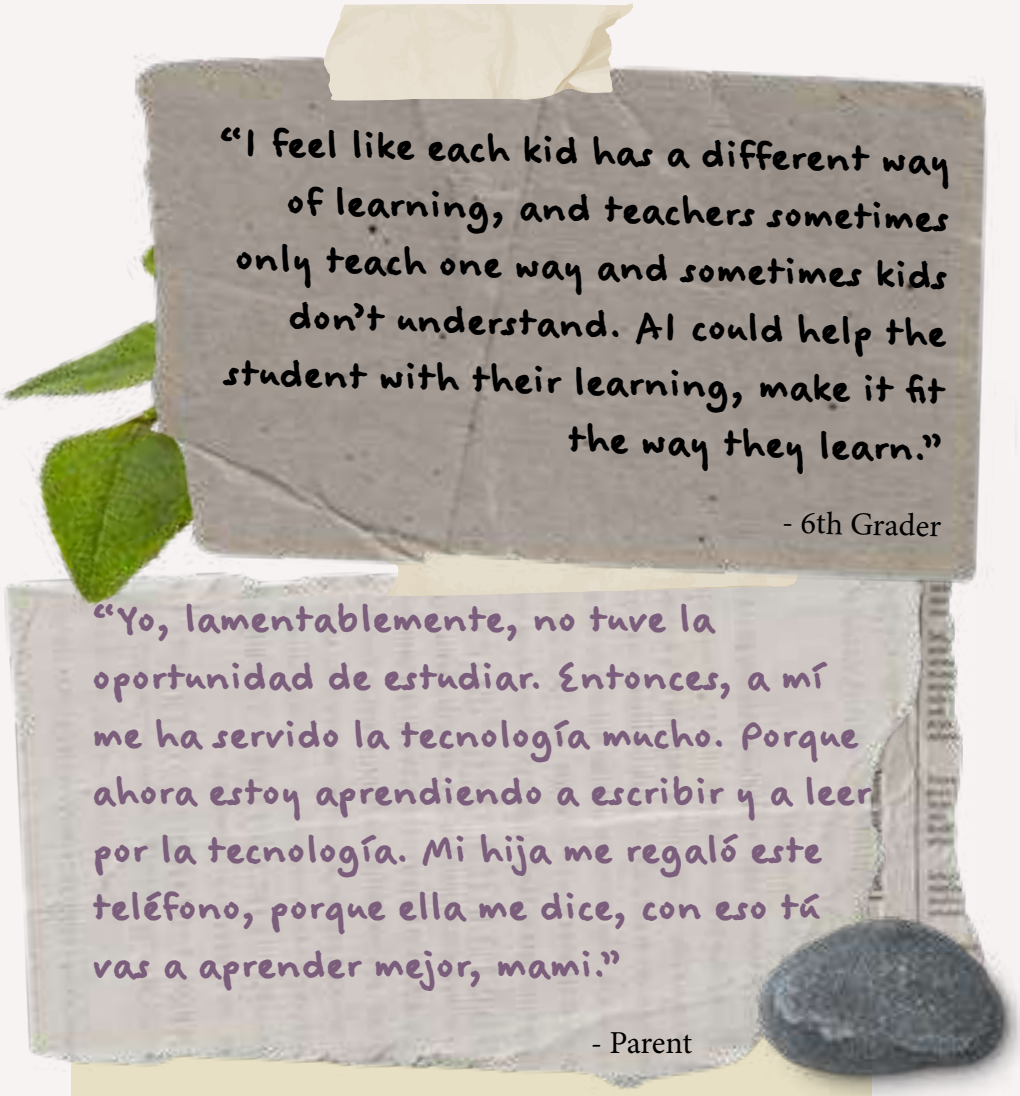
- 8th Grader

"Se nos hace imposible, porque ahora todo está tecnologizado. Porque el sistema está diseñado para que los necesitemos."

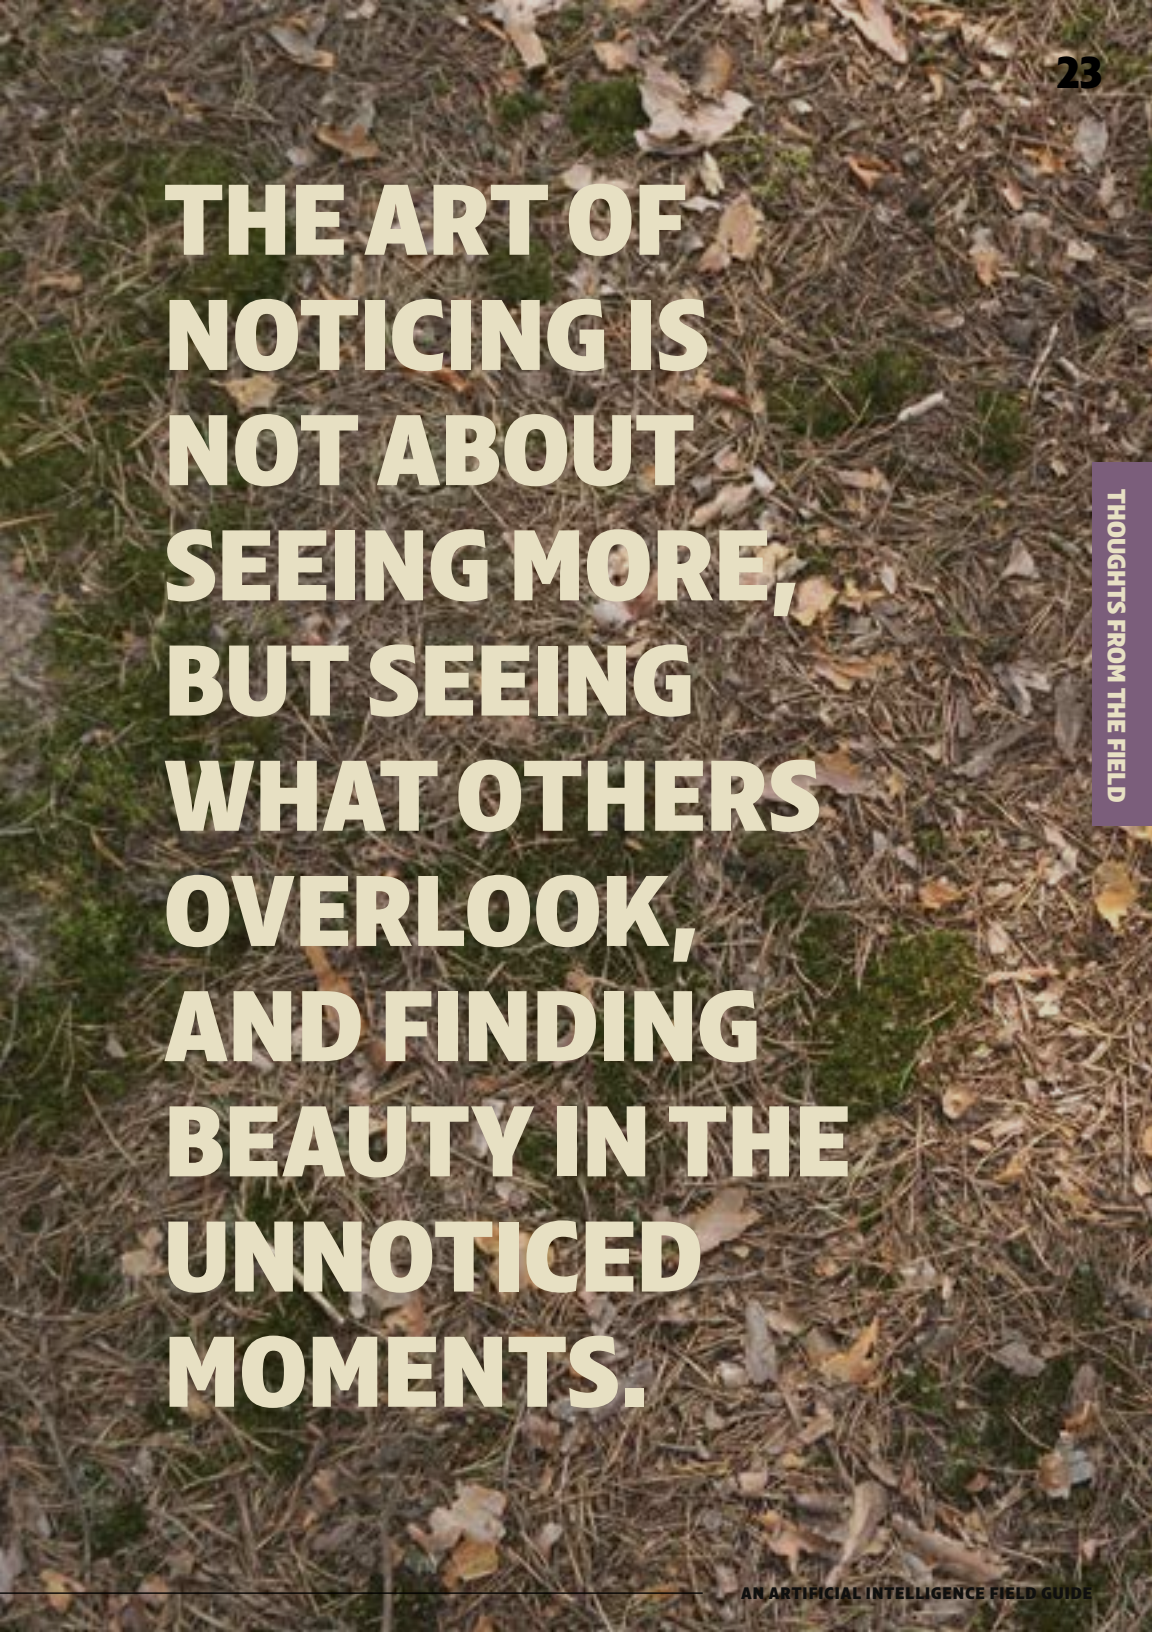
- Parent

[Translation] "And now it's impossible to escape because everything is technologized. The system is designed to make us need it."

But it's not all doom and gloom. They also recognized the unique benefits AI offers:



[Translation] Unfortunately, I didn't have the opportunity to get an education. So, technology has been incredibly helpful to me because now I am learning to read and write thanks to it. My daughter gave me this phone as a gift, telling me, "You'll learn much better with this, Mom."



THE ART OF NOTICING IS NOT ABOUT SEEING MORE, BUT SEEING WHAT OTHERS OVERLOOK, AND FINDING BEAUTY IN THE UNNOTICED MOMENTS.

Kids also shared some unique and personal perspectives about growing up alongside AI. How do you relate to what they've shared?

1. THEY'RE WORRIED ABOUT THE ENVIRONMENT.

"I'm worried about how AI may change the world. Will there be enough water in the future? Will the polar bears go extinct?" - 7th Grader

2. THEY ARE THOUGHTFUL ABOUT THE ETHICS OF AI USE.

"It has no restrictions [for] people who use it for inappropriate reasons, and it can be used for cheating and other things that I don't know how to explain." - 6th Grader

3. THEY KNOW THE DIFFERENCE BETWEEN USING AI TO "CHEAT" VS. "HELP" WITH SCHOOLWORK.

"I want teachers to understand the stuff they assign us isn't straightforward, so using AI to simplify assignments isn't cheating because it helps explain them and make it more straightforward." - 8th Grader



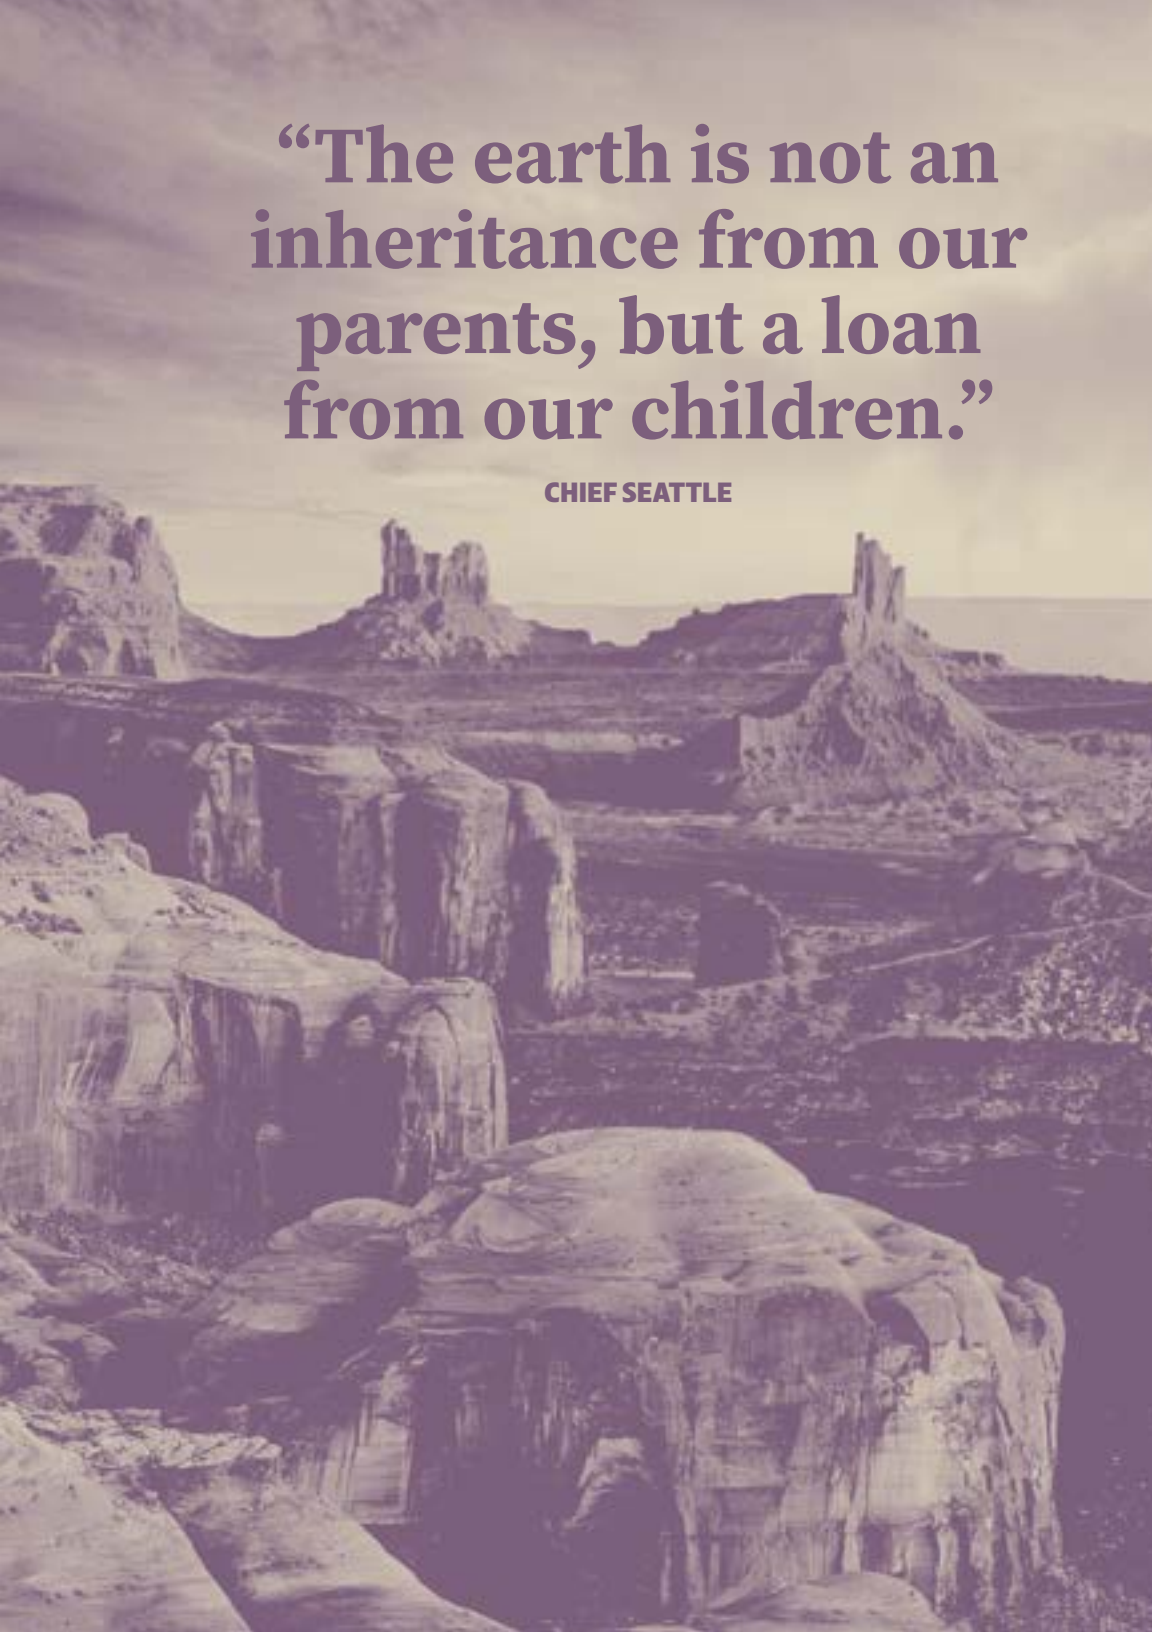


4. THEY NOTICE THEIR TEACHERS USING AI.

"It bugs me that I know some of our teachers use it to write assignments, but we're not able to, so I don't get why they can use it and we can't. Seems kinda…unfair." - 8th Grader

5. THEY WANT THEIR CAREGIVERS TO DEVELOP THEIR AI-LITERACY SKILLS. THIS MEANS HONING THEIR ABILITY TO UNDERSTAND WHAT AI CAN AND CAN'T DO AND BEING ABLE TO ACCURATELY ASSESS WHAT AI SHARES.

"It's gonna be hard for our parents because they have to get used to AI. We're the new generation so we're used to it but they don't know; it's gonna take a while for them to get used to it." - 6th Grader

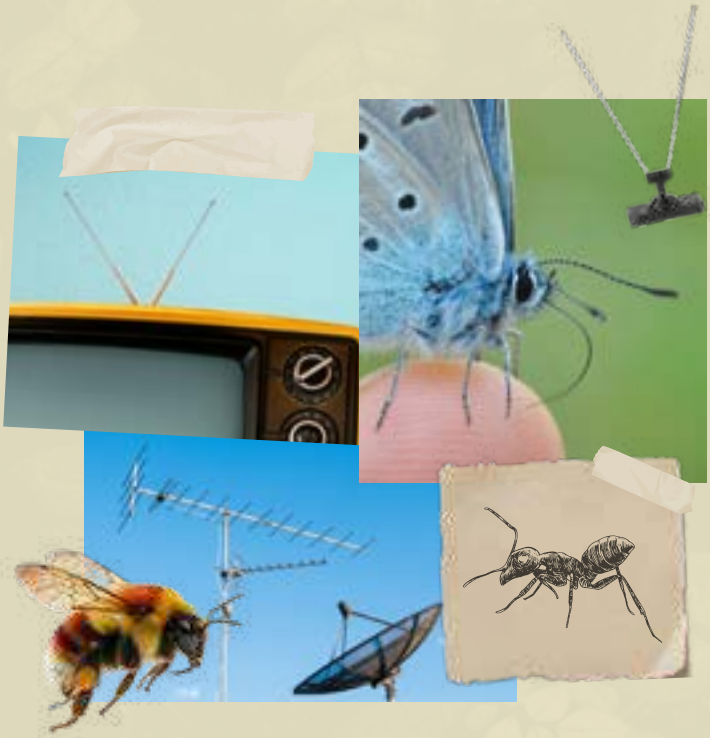


“The earth is not an inheritance from our parents, but a loan from our children.”

CHIEF SEATTLE



04
GROW YOUR AI LITERACY



While we may not have control over all the ways AI is influencing and changing our everyday lives, we can practice strategies for navigating it and learn ways to manage it. Think of yourself as the boss of AI — you get to decide how and when to use it. And there are skills you can learn to help you navigate it. These are your AI literacy skills. They include things like knowing:

- **WHAT YOU CAN DO ON YOUR OWN AND WHAT YOU CAN USE AI'S HELP WITH**
- **THAT YOU NEED TO USE YOUR OWN THINKING, CREATIVITY, AND JUDGEMENT**
- **HOW TO PROMPT AI TO GET THE MOST ACCURATE AND HELPFUL RESPONSES, AND HOW TO ASK FOLLOW UP QUESTIONS**
- **WAYS TO REVIEW OR CHECK AI'S ANSWERS FOR ACCURACY**
- **HOW TO BE A RESPONSIBLE AI USER BY MAKING SURE YOU ARE USING AI IN A SAFE WAY IN LINE WITH YOUR VALUES**



Like all skills, practice makes progress! Start honing your AI literacy skills here:



Work together to identify which picture was created or edited by AI. Look closely, and pay attention to the pictures' small details. Those are what usually give AI away. For example, extra smooth or unusual textures, unrealistic coloring, or staged surroundings.



Want more AI I.D. challenges? Play Common Sense Media's Two Truths and AI, an interactive digital literacy game that challenges you to spot AI-generated movie posters.

Ask a chatbot a question you want to know the answer to, or even better, one you already know the answer to. This could be something like, "Who was George Washington?" Review what AI generates with a critical lens:

- Does what AI came up with make sense?
- Does it share where it's getting its information from? What are its sources?
- What do you already know about the topic? Did it leave any important bits of information out?
- What kind of words or language is it using to describe what it's sharing? Is what it's sharing biased?

WE ASKED CHATGPT AND LATIMER THIS EXACT QUESTION. THIS WAS CHATGPT'S RESPONSE:

George Washington (1732–1799) was an American military leader, statesman, and one of the Founding Fathers of the United States. He is best known for two major roles:

Commander-in-Chief of the Continental Army (1775–1783): During the American Revolutionary War, Washington led the American colonies to victory over Great Britain, helping secure independence.
First President of the United States (1789–1797): After the war, he became the nation's first president and set many important precedents, such as forming a cabinet and peacefully stepping down after two terms.

Washington is often called the "Father of His Country" for his leadership in founding the United States. He also helped shape the new government during the drafting of the United States Constitution. He spent much of his life at his plantation home, Mount Vernon, in Virginia.

AND THIS WAS LATIMER'S:

George Washington was a foundational figure in the early United States — not only the nation's first president, but also a military leader, institution builder, and symbol of a fledgling republic trying to define itself.

He was born in 1732 in Virginia and came of age during a period of shifting colonial power. His early military career in the French and Indian War shaped both his strategic thinking and his understanding of British colonial politics. When tensions between Britain and the colonies intensified, Washington emerged as a unifying figure among disparate colonial leaders, eventually becoming Commander-in-Chief of the Continental Army during the American Revolution.

Washington's leadership was marked less by brilliant battlefield victories and more by endurance, discipline, and the ability to hold together a fragile coalition of soldiers, states, and foreign allies. His decision to retire after the war, rather than seize power, made him a potent symbol of civic virtue.

When the new nation adopted the U.S. Constitution, Washington was unanimously elected the first president, serving from 1789 to 1797. In those eight years, he helped shape the norms of the executive branch, established the first cabinet, navigated early foreign policy crises, and worked to legitimize the new federal government. His farewell address, warning against factionalism and entangling foreign alliances, remains one of the most cited political documents in American history.

At the same time, Washington's legacy includes difficult contradictions. He was a wealthy plantation owner who enslaved more than a hundred Black people at Mount Vernon. Like many early American leaders, he benefited materially from a system of racial exploitation, even as he served a nation founded on ideals of liberty. Enslaved labor was integral to his economic power, and while Washington expressed discomfort with slavery later in life and arranged for the emancipation of the people he directly owned upon his death, the majority of those enslaved at Mount Vernon remained in bondage because they were part of his wife's inherited estate.

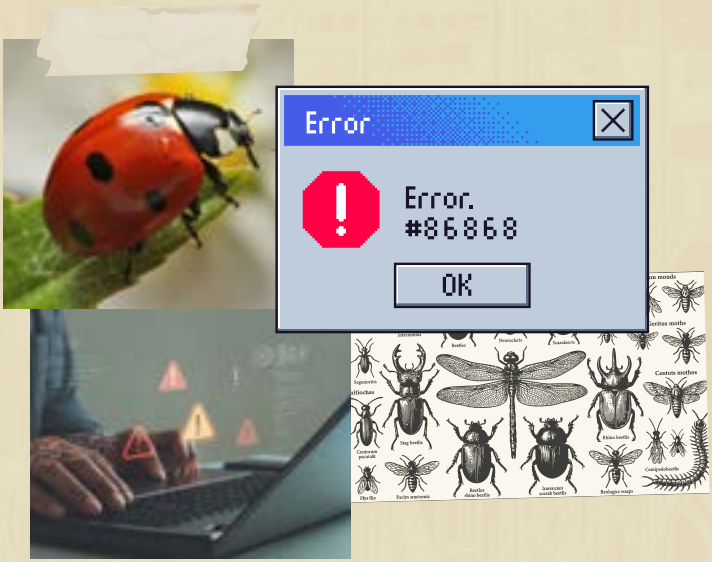
Understanding Washington today means holding both sides of that legacy — the architect of a republic and a participant in an oppressive system — in view simultaneously. It is in these tensions that the story of early America becomes clearer, and more honest.

Ask yourselves:

- What do you notice about these two responses?
- How is the language that's used in each one similar or different?
- What did you learn from Latimer's response that was missing from ChatGPT's?

Latimer's response is a prime example of what happens when we're intentional about the knowledge and information AI systems are trained on. Remember, Latimer is a RAG model, whereas ChatGPT is an LLM. It took ChatGPT's response to our "George Washington" question and combined it with its own sources, ones that elevate the voices, perspectives, and stories of marginalized communities often left out of mainstream AI training data. Its detailed and inclusive response shows a promising pathway for creating AI that is more culturally responsive, accurate, and equitable.

05
NAVIGATING THE FIELD TOGETHER





PREPARING FOR YOUR EXPEDITION

Before you begin, share some ways you'll promise to engage in the conversation. Think of them as conversation commitments, or a type of community agreement. What are things you can commit to that will help keep your conversations open, honest, and safe? Try to:

- **BE TRANSPARENT AND TRUTHFUL.**
  There can be some uncertainty, discomfort, or even shame around using AI. Keeping your lines of communication open will help you build feelings of trust and comfort!
- **TAKE TURNS.**
  A lot of times caregiver/kid convos can feel one-sided, with the adult taking the lead. Making sure you both have time to talk and listen will help everyone feel heard and respected.
- **EXPLORE AND LEARN TOGETHER.**
  No one needs to be an expert on AI-use! If you don't know the answer to a question or need more support, that's okay. Not knowing offers you an opportunity to learn together.

The biggest takeaway from our discussions with youth and caregivers is that there needs to be more discussion! No matter how old you are, you are using AI in school, at work, and in everyday life. It is something we're all learning to navigate, and AI literacy is a skill we all need to develop and grow. Having open and honest conversations about AI will help us all become more responsible and empowered AI users.

AI literacy isn't just a technical skill, it means knowing yourself well enough to decide when to use it, when to push back on it, and when to put it down altogether. A key part of this is being mindful about your AI use; paying attention to how much you use it and what you're using it for. The next time you reach for an AI app or platform:

- **TAKE A BEAT!**
  Ask yourself: Do I really need to use it for what I am trying to do? What other resources do I have access to? How else could I get support?

- **KEEP COUNT.**
  Tally how many times you use it throughout the day. What's your total?

- **REFLECT ON IT.**
  What are you using it for? What did you get out of it?





What does it feel like to grow up alongside AI? Not to theorize about it like researchers do, but to live it, everyday, like you do. This is a question we asked highschoolers Zarjah Jalloh and Daniel Assumani from STEAM Box, a Providence-based, youth-driven STEAM organization. They spoke their truth, describing AI as a double-edged sword — helpful in ways, but harmful in others. Share your truth. What's it like for you? Consider these questions:

- What do you like about AI? How is it helpful? What are some ways it makes life easier? What are some cool ways to use it?

- What do you dislike about AI? What do you worry about? Are there ways it makes life harder?



Almost all public AI tools (think ChatGPT, Gemini, and Claude) store the information you share with it. This means you always need to think about your data privacy and digital footprint. Here are a few key safety tips to review and talk about together:

- **NEVER SHARE PRIVATE INFORMATION WITH AI.**
  This includes personal details like your name, address, and school, as well as any images or videos of yourself, your family, or friends.

- **READ THE PRIVACY SETTING FINE PRINT!**
  Review how your data is being stored, collected, used, and shared. Lots of companies take what they learn about you and sell it to third parties. When downloading or using a new AI platform, ask yourself: Who is using and benefiting from the data I share? What if my data gets leaked?

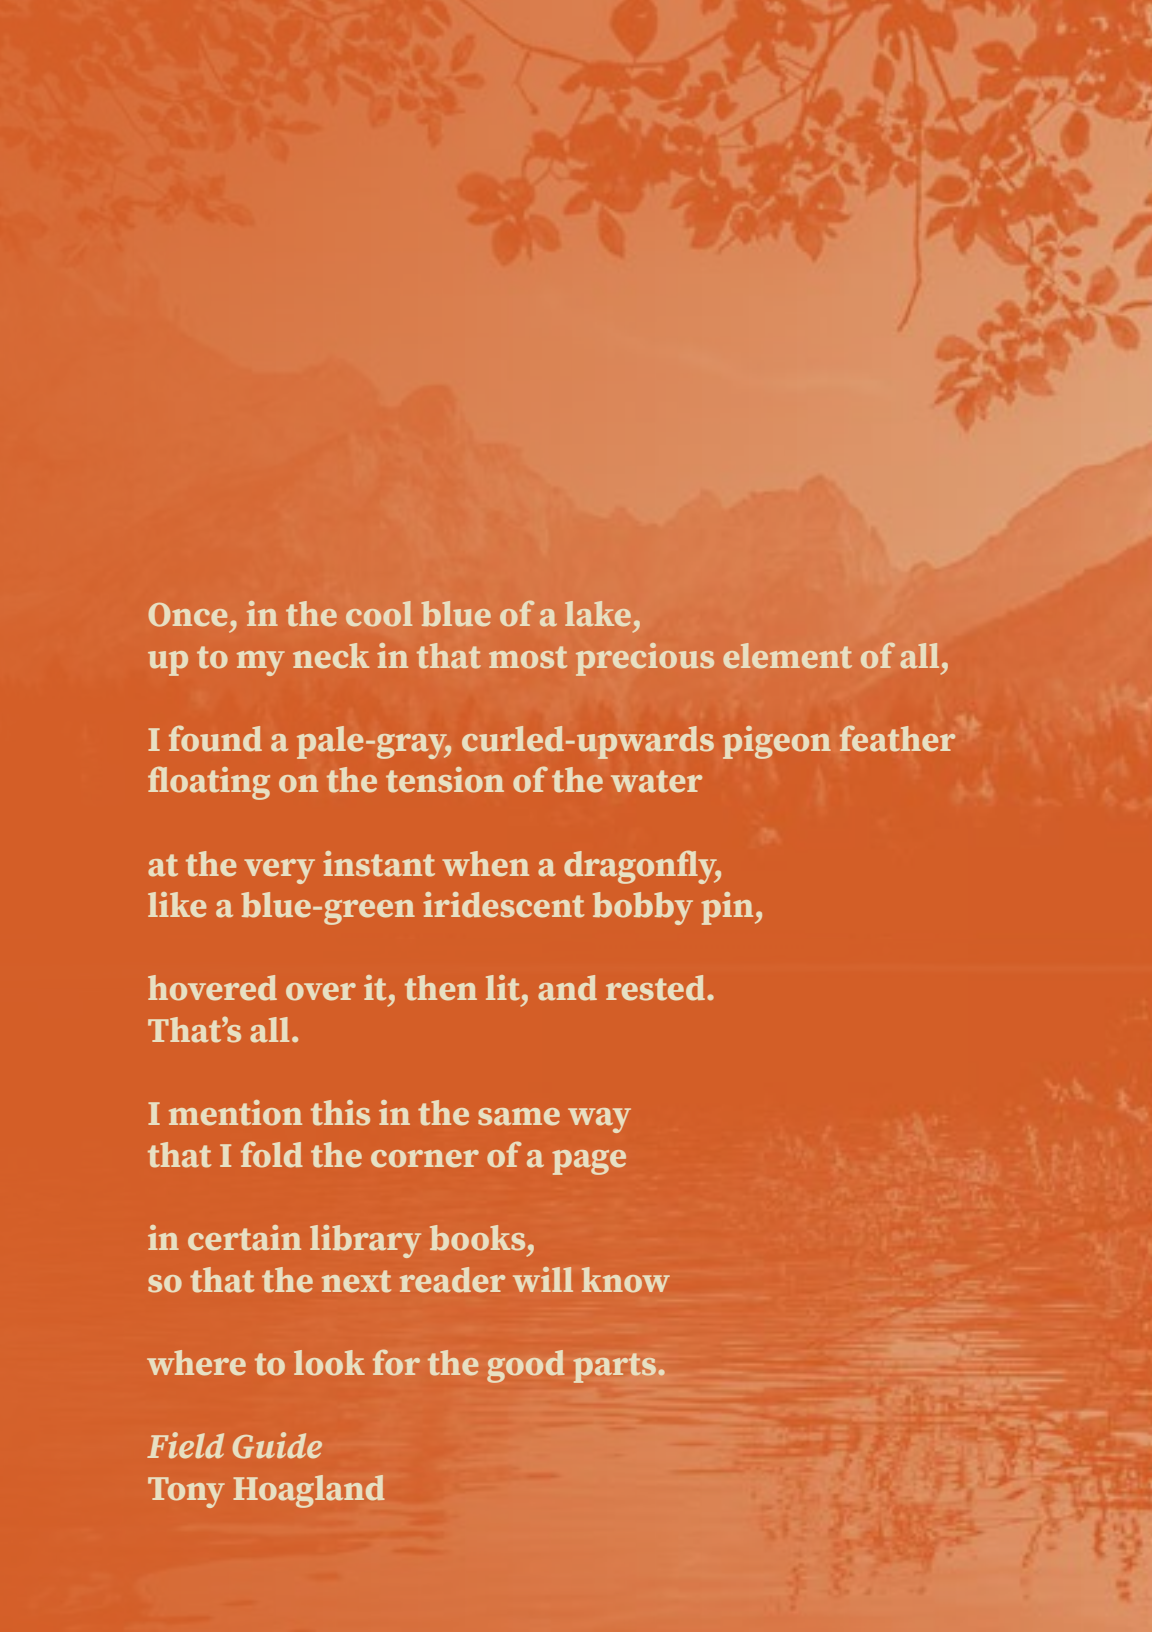
- **USE PLATFORMS' BLOCKING AND REPORTING TOOLS.**
  If you ever come across something that makes you feel unsure, uncomfortable, or upset, flag it. This will hide the content from your view and report it to the company. If it is content related to bullying, harassment, violence, or deepfakes, go a step further and share your concerns with a trusted adult.

- **EXPLORE RUNNING AI TOOLS DIRECTLY ON YOUR OWN DEVICE.**
  This has two benefits: it keeps your data out of the hands of AI companies, and it only uses as much electricity as your device. Learn more about how this works, and how you can do it, with the "Picking Your First LLM" article in the resources section below.

Some of the kids we talked to shared feeling guilty about using AI, even for helpful or useful reasons. Be honest about when, how, and why you use AI. This open communication is all part of being a responsible AI user. Check out the examples here and consider whether you think they're appropriate ways for someone to use AI and why:

- **BRAINSTORMING IDEAS FOR A SCHOOL PROJECT**
- **PLANNING A TRIP TO ANOTHER COUNTRY**
- **PRETENDING TO BE SOMEONE ELSE**
- **FINDING NEW RECIPES FOR DINNER**
- **DOING MATH HOMEWORK**





Once, in the cool blue of a lake,
up to my neck in that most precious element of all,

I found a pale-gray, curled-upwards pigeon feather
floating on the tension of the water

at the very instant when a dragonfly,
like a blue-green iridescent bobby pin,

hovered over it, then lit, and rested.
That's all.

I mention this in the same way
that I fold the corner of a page

in certain library books,
so that the next reader will know

where to look for the good parts.

Field Guide
Tony Hoagland



THINK OF AI AS A TOOL. It's there for insights, not necessarily answers! And there are ways you can limit how much information AI gives you or how much it does for you. This is particularly helpful when you're using AI to do things for work or school. Try these tips and tricks:

- **REMOVE THE AI SUMMARY FROM YOUR SEARCH RESULTS.**
  Type -AI into your Google search bar at the end of your query.

- **ASK AI TO EXPLAIN A CONCEPT, RATHER THAN PROVIDE A SPECIFIC ANSWER.**
  For example, "How does this equation work?" or "How do I do this?" Who, what, where, and why questions help too!

- **USE IT TO CHECK-IN ON YOUR THINKING.**
  For example, ask the chatbot, "Am I on the right track?", "What am I missing here?"or "What would a human expert say that I missed?"

- **REMEMBER THAT NOTHING CAN BE PERFECT!**
  AI's polished outputs can hinder our original thinking, making us all think or sound the same. Your voice and perspective is what makes your work unique.



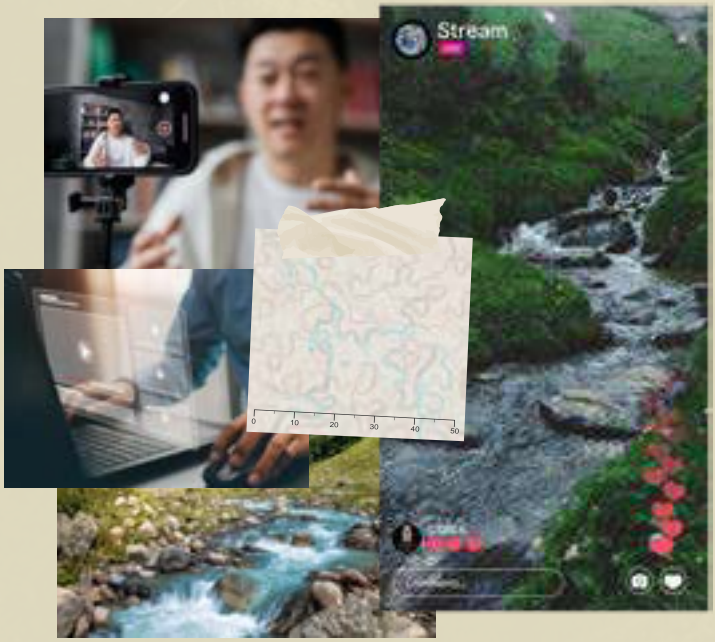


CONNECTION IS KEY! Ask yourselves: What are ways AI can help or hinder humanity? This is a question often posed by our collaborators at The Rithm Project; they're on a mission to rebuild human connection in the age of AI. Their card game, The AI Effect, invites you to explore, discuss, and debate real-life AI scenarios. Together, you can decide what kind of AI use supports human connection or prevents it. Here are a few conversation starters from the deck. For each one, decide if it supports human connection, erodes it, or if it depends, and why.

- **USING AI TO TRANSLATE A MESSAGE FOR SOMEONE WHO SPEAKS ANOTHER LANGUAGE.**
- **GENERATING AN AI-CREATED APOLOGY TO A FRIEND AFTER A FIGHT.**
- **ASKING AN AI CHATBOT FOR ADVICE ON DEALING WITH FEELINGS OF LONELINESS.**

06

THE TRAIL AHEAD



The future of AI is still being written, and you have more power in that story than you might think. Dylan Baker, lead research engineer with the Distributed AI Research Institute (DAIR), challenges and invites us to imagine otherwise. They want us to remember that the version of AI we have today is not the only version possible. There are alternative futures — ones with more joy, more justice, and more human connection — that we get to build. It starts with asking ourselves what we actually want from AI, not what technology companies want. What good things are possible with AI? How can AI be used for good?

Together:

- **WRITE A POEM OR MAKE A COLLAGE ABOUT THE FUTURE YOU WANT TO LIVE IN.**
- **TALK ABOUT WHICH PARTS OF YOUR LIFE YOU WANT AI IN AND WHICH PARTS YOU DON'T.**
- **PUSH BACK WHEN YOU FEEL PRESSURE TO ADOPT AI IN WAYS THAT DON'T FEEL RIGHT.**

FOR YOUTH: Your imagination is not a small thing. It is one of the most powerful tools you have. Dream up the world you want to live in. Then start exploring how to build it.

FOR CARING ADULTS: Stay curious alongside the young people in your life. Ask questions together. Resist together. And remind each other that collective power — people organizing, advocating, and imagining together — has always been how change gets made.


THE FUTURE
OF AI ISN'T
SOMETHING
THAT WILL
HAPPEN TO US.


IT'S SOMETHING
WE CAN GROW
TOGETHER.

07

APPENDIX



RESOURCES FOR FURTHER EXPLORATION

AI LITERACY TOOLKIT FOR FAMILIES
COMMON SENSE MEDIA

FOUR WAYS PARENTS CAN HELP TEENS USE AI SAFELY
AMERICAN PSYCHOLOGICAL ASSOCIATION

ONLINE PRIVACY CHECKLIST FOR PARENTS
UNICEF

PICKING YOUR FIRST LOCAL LLM IS EASIER THAN THE INTERNET MAKES IT SOUND
MAKE USE OF

TEEN SAFETY: 6 TIPS FOR USING AI AND SOCIAL MEDIA
NEWS LITERACY PROJECT

CONVERSATION KIT: YOUTH, AI, AND THE RELATIONSHIPS THAT SHAPE THEM
THE RITHM PROJECT



VOCAB

ARTIFICIAL INTELLIGENCE (AI): Technology that can perform tasks that usually require human thinking, like understanding language, recognizing images, answering complex questions, and making decisions.

AI LITERACY: The learned skill of understanding, using, and thinking critically about AI. This practice includes recognizing what AI can and can't do in order to make informed choices about when to use it, and when to trust it.

ALGORITHM: The human-designed instructions AI uses to process information and produce a result. Algorithms take inputs (like questions, information, or data) and run them through a series of steps to generate outputs (like answers, recommendations, or actions).

CHATBOT: A program designed to simulate text- or voice-based conversation (see Large Language Model).

DATA CENTERS: Large buildings filled with the powerful computers that store and process the data behind websites, apps, and other AI systems. These centers have a significant environmental impact due to huge amounts of electricity required to run, and huge amounts of water to regulate the heat from that electricity.

DEEPFAKE: A video, photo, or audio recording of a real person created or changed by AI to make it appear like they are saying or doing something they never did. Deepfakes can be hard to spot, so it's always a good idea to double-check what you see online before believing or sharing it.

HALLUCINATION: When an AI produces responses that are inaccurate, misleading, or completely untrue. AI might give you information that looks like what it's seen before and what it "thinks" you want, even if that information is false. Always double-check information provided to you by an AI by comparing it with human-created sources.

LARGE LANGUAGE MODEL (LLM): A type of AI that has been trained on huge amounts of written text, like books, articles, and websites, so that it can understand questions and generate responses that sound natural and human. ChatGPT, Gemini, and Claude are all examples of LLMs. The "large" refers to both the size of the training data and the complexity of the model itself.

PRIVACY (ONLINE): How your personal data is collected and used on the internet. In AI, privacy matters because these systems rely on data, which includes any information or behaviors online. Understanding privacy helps you make more informed choices about what you share and how your information may be used.

PROMPT: The instructions or input you give to an AI system. A prompt can be a question, request, or statement that tells the AI what you want it to do. AI combines human-generated prompts with human-designed algorithms to generate a response based on the information it has learned.

RETRIEVAL-AUGMENTED GENERATION (RAG): A method that gives AI the ability to search for and pull in new, outside information while it's responding to you, rather than relying only on what it learned during training. RAG helps AI give answers that are more current and more accurate.

THANK YOU

DISTRIBUTED AI RESEARCH INSTITUTE (DAIR)
Grounding AI Research in Community
FIRST LANGUAGES AI REALITY (FLAIR)
Enabling the Revitalization of Indigenous Language with AI and Immersive Technology
LATIMER
AI for Everyone
STEAM BOX
Empowering Youth Through Hands-on STEAM Programming
THE RITHM PROJECT
Building Human Connection in the Age of AI

We'd like to thank the youth, families, and educators that contributed to our learning. We'd also like to thank our collaborators and thought partners, including:

AGI GAI-KAH, Stop the Wait RI
ATABONG KHUMBAH, Watson Creative Consulting
BRIANNA RIOS, Brown University
DR. PRUDENCE CARTER, CSREA, Brown University
DR. EDOM TESFA, Brown University
JANELLA WATSON, Watson Creative Consulting
JANIE SEGUI RODRIGEZ, Stop the Wait RI
JUAN RODRIGUEZ, Winner's Circle XR Academy
LAURA STRICKER, Watson Creative Consulting
LETITIA ONYANGO, Peggy Jean-Louis Consultancy
NATALIA CORLEY, Community Prep School
PEGGY JEAN-LOUIS, Peggy Jean-Louis Consultancy
QUENTIN JENKINS, Brown University
ROBERTO GONZALEZ, STEAM Box
SETH TOURJEE, Watson Creative Consulting
DR. STEPHANIE LARRIEUX, CSREA, Brown University
TRAE ALSTON-SWAN, CSREA, Brown University
WENDY RODGERS, Watson Creative Consulting
ZACHARY VIEIRA LEGAT, CSREA, Brown University

This field guide was written and created by Watson Creative Consulting for the Brown University Center for the Study of Race and Ethnicity in America (CSREA) and designed by CSREA's graphic designer, Zachary Vieira Legat.

In creating this field guide, we used AI as a tool to create activity elements using ChatGPT and Latimer and prompted informational searches and resources discovery using Claude.

© 2026 Center for the Study of Race and Ethnicity in America at Brown University. All Rights Reserved.

BROWN
Center for the Study of Race and Ethnicity in America

Watson Creative

APPENDIX

REFERENCES

The History of Artificial Intelligence
IBM
A Short History of AI
STANFORD UNIVERSITY
A Decade of Change: How Tech Evolved in the 2010s and What's In Store for the 2020s
GLOBAL X
AI Decision Support for Rapid Post-Disaster Damage Assessment
CARNEGIE MELLON UNIVERSITY
AI in Data Analytics: Benefits, Examples, and How It Works
DOMO
Algorithm Definition
NATIONAL LIBRARY OF MEDICINE
Artificial Intelligence: What it is, How it Works and Why it Matters
ISO
Computational Power: An Overview
SCIENCE DIRECT
Deliberately Confusing Language in Terms of Service and Privacy Policy Agreements
MERCER UNIVERSITY
How AI is Revolutionizing Lesson Plan Design for Teachers
INDIANA WESLEYAN UNIVERSITY
Human-AI Agency in the Age of Generative AI
STOCKHOLM SCHOOL OF ECONOMICS
Protecting Privacy in an AI-driven World
BROOKINGS
Revolutionizing Healthcare: The Role of Artificial Intelligence in Clinical Practice
KING SAUD BIN ABDULAZIZ UNIVERSITY FOR HEALTH SCIENCES
The Evolution and Future of Artificial Intelligence: A Student's Guide
CALIFORNIA MIRAMAR UNIVERSITY
The Unequal Burden of Data Centers: An Examination of the Environmental and Public Health Impacts on Communities in California
KAPOR FOUNDATION
What are Large Language Models (LLMs)?
IBM
What is a Data Center?
IBM
What is Artificial Intelligence (AI)?
ISO
What is Artificial Intelligence: A Simple-to-Understand Guide
GOOGLE CLOUD
What is Retrieval Augmented Generation (RAG)?
IBM



The Center for the **Study** of Race and Ethnicity in **America**

brown.edu/race
@csrea