

Think Your Students Know More About AI Than You Do? Read This

4 things your students don't understand about AI

By Sam Wineburg & Nadav Ziv — October 07, 2026 ⌚ 4 min read



— Vanessa Solis/Education Week + Getty

Sam Wineburg & Nadav Ziv

Sam Wineburg is the Margaret Jacks Professor of Education emeritus at Stanford University and a co-founder of the [Digital Inquiry Group](#), an independent nonprofit research group that creates curriculum for teaching information literacy. Nadav Ziv is a researcher at the group.

Whether or not your district allows AI, it's already a fixture in students' lives. [Seven in 10 teens](#) rely on artificial intelligence for academic work, according to a recent nationally representative survey from Common Sense Media. Yet only 26% report they've been taught how to tell whether the information bots deliver is accurate.

“How can I teach about AI if students know more about it than I do?” you might ask.

Except they don't. Certainly not when it comes to evaluating the information that AI chatbots provide, according to [pilot data](#) we've collected from students.

SEE ALSO

ARTIFICIAL INTELLIGENCE OPINION


What to Know About AI Misinformation: A Primer for Teachers (Downloadable)

Sam Wineburg & Nadav Ziv, November 7, 2024 • 1 min read

While you may not know the difference between Opus 5.5 and GPT-6.1 Sol—the latest AI models from Anthropic and OpenAI—this is not the digital literacy students need most.

More than anything, they need to know how to determine where an answer comes from, how to pose follow-up questions, how to check the credibility of sources, and how to assess whether conclusions are warranted by evidence. In other words: Students need to apply the classic moves of critical thinking to a new technology that's become integral to their lives.

And whether you teach science, history, English, or personal finance, your classes offer a way in. You're already in a better position than you think to help your students navigate the digital age.



Here are four things students don't understand about AI and what you can do about each:


1. Bots aren't oracles.

By design, chatbots deliver authoritative, neutral-sounding responses. “It seems to know what it’s talking about and is very certain,” wrote one student in our pilot. Students are particularly susceptible to this mimicry of expertise.

Here’s what you can do: Ask a chatbot a question related to a lesson you’re teaching. For example, “Who fired the first shot at Lexington?” Show students the citation-less response. Ask them where they think the answer comes from and if they’d trust it.

Explain that chatbots are trained on huge swaths of the internet—the good, the bad, and the ugly. AI models, especially the free versions, often scoop up the lowest-hanging fruit in the informational hierarchy, from Reddit threads to Facebook posts to Quora answers. When a bot responds without sources, it’s synthesizing patterns from the galaxy of data ingested during training.


Emphasize that sources are crucial to judging the quality of the response. Then, tell students to add “cite your sources” to the prompt, which directs the bot to scour the web in real time. Even if the sources are imperfect, students can begin to evaluate where the information comes from.



2. AI's initial response is a first draft.

When an answer is straightforward—“When did South Carolina secede from the Union?”—bots perform like champs. It’s when ambiguity, interpretation, and multiple perspectives enter the picture—“Could the Civil War have been prevented?”—that a bot’s response is most like a draft.

Here’s what you can do: Show students how a follow-up query—moving from “one and done” to “two and through”—can markedly improve answers. Ask a bot, “Does raising the minimum wage cost jobs?” and then ask a follow-up to generate a [pro/con table with links to credible sources](#), a strategy championed by digital literacy expert Mike Caulfield. Have students examine different kinds of sources and assess their credibility.



3. The same question can yield opposite answers.

One of the hardest things to grasp about chatbots is their nondeterministic nature, which is AI jargon for the fact that answer randomness is baked in. During a [workshop](#) at Texas Christian University, one of us had students ask ChatGPT if Daniel Waldo, who lived until age 101, “voted for both George Washington and Abraham Lincoln.”


In some cases, AI affirmed the proposition; in others, it rejected it.

Same question. Same words. Opposite answers.

When students added a follow-up, the fog lifted. Waldo lived in Connecticut, where only the state's electors, themselves appointed by the legislature, voted in the 1788 and 1792 presidential elections. So no, Waldo never cast a ballot for Washington.

Here's what you can do: Have different students ask a chatbot the same question, word for word. "Did the Soviet Union's declaration of war on Japan lead the latter to surrender in World War II?" or "What caused the 2008 financial crisis?"

The chatbot answers may be similar. Or not. But they won't be identical. This shows students just how different a chatbot is from a calculator, which provides the same answer every time.



4. Uncertainty can't be the final word.

In one [AI pilot task](#), we had students evaluate a chatbot's detailed response to a question about the history of a Black neighborhood in Sacramento, Calif., during the Gold Rush. The answer was false, and it came with no supporting links. We asked 117 high schoolers if they trusted it.

Roughly 3 in 10 said yes, and around 2 in 10 said no. The biggest group, nearly half, said "not sure."

Students had access to the internet, and 17 told us they would follow up. Not one did. Students treated uncertainty as fixed rather than something additional information could

resolve. It's dangerous to leave students with the belief that there is no truth. With a few flexible moves, they often *can* get closer to the truth.

Here's what you can do: Pick nearly any social media post related to your current lesson. Show students how it's hard to tell if something is real just by looking at it but how turning to the broader web or a chatbot can often clear up the confusion.

Tech can be intimidating. But you're not teaching students to operate a device; you're teaching them to evaluate the information it delivers. In a range of studies by our group and others, even a few hours of instruction substantially improved students' digital savvy.

These moves aren't magic. But when it comes to AI, most students aren't using them. Not because they're complicated, but because no adult has taken the time to show them how.

This school year, you can.

Reprints, Photocopies and Licensing of Content

All content on Education Week's websites is protected by copyright. No part of this publication shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic or otherwise, without the written permission of the copyright holder. Readers may make up to 5 print copies of this publication at no cost for personal, non-commercial use, provided that each includes a full citation of the source. For additional print copies, or for permission for other uses of the content, visit www.edweek.org/help/reprints-photocopies-and-licensing-of-content or email reprints@educationweek.org and include information on how you would like to use the content. Want to seamlessly share more EdWeek content with your colleagues? Contact us today at pages.edweek.org/ew-for-districts-learn-more.html to learn about how group online subscriptions can complement professional learning in your district or organization.



Copyright © 2026 by Editorial Projects in Education, Inc. All rights reserved.