

Generative AI in K-12 Education: Challenges and Opportunities

Updated August 2024

Generative AI in K-12 Education: Challenges and Opportunities

Updated August 2024

Table of Contents

Introduction

- How to read this paper
- Background
- Responsible AI
- Education-specific considerations
- Examples of principles and guidelines for responsible AI
- Cautionary notes

How Americans view AI

- General attitudes
- Awareness and knowledge
- Usage

AI literacy among teachers

The state of school policy

Policy approaches in the U.S., across the Atlantic, and beyond

- United States
- European Union
- United Kingdom and Australia
- Governance models

Opportunities

Student-supporting

- Adaptive learning
- Creativity
- Project-based learning and collaboration

Teacher-supporting

- Lesson planning and content differentiation
- Coaching and feedback
- Grading and assessments

System-supporting

- Data interoperability
- Administration and logistics
- Parent engagement

Seckinger High: All In on AI

Risks, harms, and challenges

- Inaccuracies and hallucinations
- Bias and inequity
- Lack of representation
- Information censorship
- Cheating and plagiarism
- Overreliance on technology and loss of critical thinking
- Data security
- Consent and privacy
- Human replacement

What students think

Recommendations

Conclusion

Acknowledgments

APPENDIX

- Resources, links, and information
- Educator snapshot: how one social studies teacher is navigating AI
- Educator snapshot: how one fourth grade teacher uses Kyron Learning
- Educator snapshot: how one fourth grade math teacher uses TeachFX

About

Introduction

Less than two years ago, generative AI moved from a developing technology still mostly confined to research purposes and academic debate to widespread usage. OpenAI publicly released their text-to-image model DALL-E 2 in September 2022,¹ followed by their AI text chatbot ChatGPT that November.² By January 2023, ChatGPT had clocked 100 million users—the fastest-growing consumer application ever.³ A year later, users could choose from several chatbot competitors, including Google Gemini and Anthropic's Claude, and thousands of applications built on these tools had flooded the market.

According to countless headlines, the advent of generative AI marked the beginning of a "revolution"⁴ across every imaginable sector of business and society, including banking, manufacturing, health care, law, entertainment, retail, design, and, of course, education. But headlines are headlines, and we've been through technological revolutions before (think automobiles, television, nuclear power, computing, and the internet in the last century, to say nothing of industry, railways, and electricity in the more distant past). What makes AI different? Is this revolution the mother of all technological revolutions? How do we keep up with, and make sense of, the latest developments?

In researching and reporting this paper, we spoke to more than 40 people about AI, running the gamut from expert PhDs to eighth graders. We also drew on the collective wisdom of Common Sense, which has built a wellspring of expertise and knowledge on generative AI over the last few years across its education, policy, research, advocacy, and growth teams. We've tried to bring together people who don't normally interact with each other in order to stitch disparate viewpoints into more of a quilt. Across that wide spectrum of voices, we heard a similarly wide range of reactions—fear, consternation, excitement, anxiety, and optimism—sometimes simultaneously. None of those reactions are wrong. Almost two years after ChatGPT's release, we lack clear guidelines, guardrails, or policy, even as we entertain great hopes for the future.

Our aim is to provide an overview of generative AI in K–12 education, including background on the technology; the state of AI literacy and school guidelines; opportunities; risks; and recommendations for the field.

¹ OpenAI. (2022, September 28). *DALL-E now available without waitlist*.

<https://openai.com/blog/dall-e-now-available-without-waitlist>

² OpenAI. (2022, November 30). *Introducing ChatGPT*. <https://openai.com/blog/chatgpt>

³ Hu, K. (2023, February 2). *ChatGPT sets record for fastest-growing user base - analyst note*. Reuters.

<https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>

⁴ Thorbecke, C. (2023, November 30). *A year after ChatGPT's release, the AI revolution is just beginning*. CNN.

<https://www.cnn.com/2023/11/30/tech/chatgpt-openai-revolution-one-year/index.html>; Candelon, F., Burtsev, M., Jha, G., Sack, D.,

Zhukov, L., & Zuluaga Martínez, D. (2023, December 1). *ChatGPT's AI revolution: How management can adapt*. *Fortune*.

<https://fortune.com/2023/12/01/chatgpt-ai-revolution-management-bcg/>; Rosenbaum, E. *MIT data show 'industrial revolution-level' leap for workers using AI*. (2023, May 10). CNBC.

<https://www.cnbc.com/2023/05/10/mit-data-show-industrial-revolution-level-leap-for-workers-using-ai.html>; Kessler, S. (2023, June

10). *The A.I. revolution will change work. Nobody agrees how*. *New York Times*.

<https://www.nytimes.com/2023/06/10/business/ai-jobs-work.html>; Fernandez, B. (2023, August 28). *Generative AI promises an economic revolution. Managing the disruption will be crucial*. *Wall Street Journal*.

<https://www.wsj.com/articles/generative-ai-promises-an-economic-revolution-managing-the-disruption-will-be-crucial-b1c0f054>

How to read this paper

Readers will note that the title of this paper focuses on "generative AI," but we discuss AI more broadly throughout; it's nearly impossible to disentangle one from the other because they've become synonymous in most people's minds. When we use the term "AI," we are referencing AI more broadly. We use "generative" and other qualifiers when that is the area of focus. Finally, given the rapid pace of change, we acknowledge that this paper represents a snapshot in time, but we've tried to focus as much as possible on considerations that won't get easily resolved in the near term.

We begin our paper with background information about AI in general, given the importance of AI literacy. We include a summary of the tenets of responsible AI, as well as examples of principles that follow those tenets, to help set a foundation for how we might approach AI in education. We also wanted to bust some commonly held myths, so we included a few cautionary notes to complement this introductory material.

Next, we run through results from multiple surveys to provide a high-level view of sentiment, usage, and understanding of generative AI among Americans in general, and young people, parents, and teachers in particular. And while survey results are informative, we also wanted to talk directly with people to get a qualitative sense of how they feel about this technology. You'll note quotes from teachers and parents throughout this paper, as well as a section dedicated to takeaways from conversations with young people. We included full Q&As with three teachers in the Appendix.

In terms of policy, we offer an overview both of select current government policy and of school/district policy. (Spoiler alert: There's much more to be done on both fronts.)

All of that—background and definitions, example principles, survey results, and policy overviews—serves as a foundation for understanding our sections on opportunities and risks. And throughout, certain themes come up repeatedly:

- Hype vs. reality
- Data security and privacy
- Transparency, trust, and accountability
- Bias and equity
- Human oversight

These same themes also run throughout our recommendations.

This is still early days for generative AI, and we don't claim to have all the answers. But we hope this paper can function as a guide to this evolving landscape and help educational stakeholders make informed decisions about when, how, and whether to use AI tools in educational settings.

Background

Ask 10 people the definition of artificial intelligence and you'll get 20 answers. Merriam-Webster defines AI as "the capability of computer systems or algorithms to imitate intelligent human behavior."⁵ That sums it up succinctly, but other, more technically accurate definitions refrain from any reference to humans, given the risks of anthropomorphizing the technology. The European Commission's High-Level Expert Group on Artificial Intelligence, for example, defines AI as "systems that display intelligent behavior by analyzing their environment and taking actions—with some degree of autonomy—to achieve specific goals."⁶ And the Google Developers glossary uses this definition: "a non-human program or model that can solve sophisticated tasks."⁷ AI has been around for decades, from recommendation engines on Netflix and voice assistants such as Siri to customer service chatbots, search engines, and monitoring systems that predict when factory machines will require maintenance. This earlier form of AI performs discrete tasks and responds to a particular set of inputs.⁸ These systems can learn from data and make predictions based on that data, but they're not producing original content.

Generative AI, on the other hand, can generate content, such as text, images, audio, and video, that is complex, coherent, and original, in response to a submitted prompt or query, by learning from large reference sets of existing data.⁹ Alexa setting a timer for your turkey while your hands are covered in raw poultry goop is convenient, but ChatGPT generating an original sonnet about your second cousin in the style of Petrarch is pretty mind-blowing.

The spectrum of reactions to generative AI we encountered as we researched this report are similar to the spectrum in the AI industry itself, though not as polarized. On one end, "doomers" advocate for a slower approach to AI development due to concerns about the speculative risks that some believe it poses to humankind's future survival.¹⁰ On the other end, "boomers" believe the positive benefits far outweigh any risks, and that development can't and shouldn't be curbed. But while the extremes get most of the attention, neither addresses the complicated reality and myriad concerns of current usage. In the middle, there's a growing effort to develop responsible, trustworthy AI that works for all people in the here and now. The vast majority of people in general also lie somewhere in the middle in how they're thinking about AI in their own lives: skeptical, but curious. And they're looking for guidance.

There's one big problem with the emphasis on preparing for a possible future world where AI rules supreme: That vision obscures present-day challenges that need to be addressed. Longtermism—the focus on ensuring that humankind will survive millions of years from now¹¹—is terrible for AI, says Tracy Pizzo Frey, a former Google AI

⁵ Merriam-Webster. (n.d.). Artificial intelligence. In *Merriam-Webster.com dictionary*.

<https://www.merriam-webster.com/dictionary/artificial%20intelligence#dictionary-entry-1>

⁶ High-Level Expert Group on Artificial Intelligence. (2018, December 18). *A definition of AI: Main capabilities and scientific disciplines*. European Commission. https://ec.europa.eu/futurium/en/system/files/ged/ai_hleg_definition_of_ai_18_december_1.pdf

⁷ Google for Developers. (n.d.) *Machine learning glossary*. <https://developers.google.com/machine-learning/glossary>

⁸ Marr, B. (2023, July 24). The difference between generative AI and traditional AI: An easy explanation for anyone. *Forbes*.

<https://www.forbes.com/sites/bernardmarr/2023/07/24/the-difference-between-generative-ai-and-traditional-ai-an-easy-explanation-for-anyone/>

⁹ Common Sense Media. *AI rating: ChatGPT*. <https://www.common Sense Media.org/ai-ratings/chatgpt>; McKinsey & Company. (n.d.). *What is generative AI?* <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-generative-ai>

¹⁰ The Sam Altman drama points to a deeper split in the tech world. (2023, November 19). *The Economist*.

<https://www.economist.com/business/2023/11/19/the-sam-altman-drama-points-to-a-deeper-split-in-the-tech-world>; Edinburgh Futures Institute. (2023, November 23). *OpenAI chief executive's firing, rehiring casts new light on the war between AI doomers, boomers*. <https://efi.ed.ac.uk/openai-chief-executives-firing-rehiring-casts-new-light-on-the-war-between-ai-doomers-boomers/>

¹¹ Samuel, S. (2023, September 7). *Silicon Valley's vision for AI? It's religion, repackaged*. Future Perfect. Vox.

<https://www.vox.com/the-highlight/23779413/silicon-valleys-ai-religion-transhumanism-longtermism-ea>

executive who led the development of Common Sense Media's AI principles. "Some people are talking about colonizing other planets and not thinking about climate change and systemic racism. It's a sleight of hand to talk about existential risk—it covers up what is happening now."

In education, Pizzo Frey sees similarities to the rise of edtech over the past 20 years, and its acceleration during and since the pandemic. "What you have at a high level is a pretty bifurcated experience: people who are really, really excited about AI and new technology and want to implement as fast as possible, and people who are really skeptical, especially frontline educators who were already skeptical of technology in education and are also skeptical of AI." AI likely won't be an overnight revolution, but more of an evolution, a series of smaller steps and changes that can add up to big transformations over time. "We could let that happen by default or we could be intentional about it. We need to question who makes decisions. Right now it's happening more by default," says Pizzo Frey.

Responsible AI

The responsible AI movement responds to that call for intentionality. Responsible AI refers to artificial intelligence that is designed, developed, deployed, and used in a way that aligns with human values and minimizes risks and harms.¹² Think of it as a focus on how to make these tools and technologies work better—that is, more accurately, fairly, and ethically—today, not in 10 years. While there's no fixed definition of exactly what responsible AI looks like in practice, consensus is emerging around certain themes.

"Hundreds of principles and charters exist across the globe, and they're remarkably similar in some predictable ways," says Pizzo Frey. "They tend to all touch on unfair bias, transparency and explainability, and privacy and safety." Here, *explainability* refers to improving human understanding of how AI works, which also builds trust in the system's results.¹³

Most people won't be able to understand the deeper technical aspects of how ChatGPT works, but everyone should be able to understand, in plain language, the basics of how the technology functions.

Most people won't be able to understand the deeper technical aspects of how ChatGPT works, but everyone should be able to understand, in plain language, the basics of how the technology functions. That also means that tech companies can't hide behind "black box" terminology as an excuse for a lack of explanation.

"Human in the loop" is a phrase that also crops up in many sets of AI principles: the need for a human to be present at certain points in the system, rather than the use of AI as a wholesale alternative to human decision-making. In education, this idea of human involvement is especially important. The U.S. Department of

¹² Nguyen, A., Ngo, H.N., Hong, Y., Dang, B., & Nguyen, B.T. (2023). Ethical principles for artificial intelligence in education. *Education and Information Technologies*. 28(4): 4221–4241. <https://doi.org/10.1007%2Fs10639-022-11316-w>; IEEE Standards Association. (n.d.). *General principles*. https://standards.ieee.org/wp-content/uploads/import/documents/other/ead_general_principles.pdf

¹³ Turri, V. (2022, January 17). *What is explainable AI?* Software Engineering Institute blog, Carnegie Mellon University. <https://insights.sei.cmu.edu/blog/what-is-explainable-ai/>

Education's May 2023 report ("Artificial Intelligence and the Future of Teaching and Learning") issues the directive to "always center educators in instructional loops," further detailing those loops as:

1. The loop in which teachers make moment-to-moment decisions as they do the immediate work of teaching.
2. The loop in which teachers prepare for, plan, and reflect on teaching, which includes professional development.
3. The loop in which teachers participate in decisions about the design of AI-enabled technologies, participate in selecting the technologies, and shape the evaluation of technologies—thus setting a context for not only their own classroom but those of fellow teachers as well.¹⁴

The 2024 National Education Technology Plan also makes a point to emphasize humans in the loop: "Regardless of technologists' claims, no emerging technology will in and of itself solve the long-standing challenges faced by educational systems. As new technologies emerge, educators must be involved in designing and developing digital tools for classroom use. Evaluating these technologies for potential classroom use at the state, school system, or school level should be iterative and include educator and student input."¹⁵ Similarly, emerging state-level guidance also places humans at the center. According to Washington state's guidance, for example: "In K–12 education, uses of AI should always start with human inquiry and always end with human reflection, human insight, and human empowerment."¹⁶

Pizzo Frey goes a step further than the "human in the loop" principle. What's needed, she says, especially in cases where AI can have a direct and meaningful impact on people's lives and opportunities (such as education), is meaningful human control. "You cannot, and from a principled perspective should not, separate humans from AI," she says. In fact, AI is sociotechnical; the technology cannot be divorced from the humans who inform and shape its use.¹⁷

Education-specific considerations

As stakeholders debate and clarify the appropriate place of generative AI in education, it's important to keep in mind such themes as bias, transparency, and safety. It's also worth highlighting some additional considerations around AI in an educational context.

First, we're talking about children. This may be stating the obvious, but the fact that children deserve greater and different protections than adults doesn't get foregrounded enough in conversations around AI in education. In a 2021 paper summarizing AI researchers' responses to questions about the ethics of AI in education, the focus on young people was associated with issues of consent, learning agency, data privacy, and algorithmic bias, as well as longer-term "impacts of engaging with young people during a formative period." Said one respondent: "We are

¹⁴ U.S. Department of Education, Office of Educational Technology. (2023, May). *Artificial intelligence and the future of teaching and learning: Insights and recommendations*. <https://www2.ed.gov/documents/ai-report/ai-report.pdf>

¹⁵ U.S. Department of Education, Office of Educational Technology. (2024.) *National educational technology plan*.

¹⁶ Washington Office of Superintendent of Public Instruction. (2024, January 18). *Human-centered AI guidance for K–12 public schools*. <https://ospi.k12.wa.us/sites/default/files/2024-01/human-centered-ai-guidance-k-12-public-schools.pdf>

¹⁷ Common Sense Media. *AI initiative*. <https://www.common Sense Media.org/ai>; Weidinger, L., Rauh, M., Marchal, N., Manzini, A., Hendricks, L.A., Mateos-Garcia, J., Bergman, S., Kay, J., Griffin, C., Bariach, B., Gabriel, I., Rieser, V., & Isaac, W. (2023, October 31). *Sociotechnical safety evaluation of generative AI systems*. Google DeepMind. <https://doi.org/10.48550/arXiv.2310.11986>

influencing impressionable young people, and that comes with a moral obligation to be as correct and proper as possible."¹⁸

Two other considerations are important to recognize as schools evaluate, review, and possibly adopt new AI-powered tools. One is that the impacts of AI in educational settings, whether via more traditional algorithms or more recent generative AI technology, aren't necessarily readily apparent. A malfunctioning AI system in a self-driving car could result in immediate, significant bodily harm or death, whereas the effects of AI on children may be harder, and take longer, to identify. Educators and school leaders nonetheless need to consider those potential downstream effects now, and research will be required to study outcomes.

The second is that AI educational tools may embody assumptions about which instructional strategies are best for learners, and not enough attention is always paid to which pedagogy is driving a specific AI tool.¹⁹ In the rush to release and evaluate new products, often there is not enough focus on the pedagogical approach that each product has adopted—or the efficacy of the tool at all. "Pedagogy-driven AI in education is often considered as good regardless of [whether] the pedagogy chosen to drive AI is an appropriate one or not in the first place," another researcher commented in that same paper on the ethics of AI in education.²⁰

Examples of principles and guidelines for responsible AI

Frameworks need road maps that spell out how they work in practice, and a number of entities have begun releasing actionable AI principles and guidelines to that end. In 2023, Common Sense Media released its first ratings of AI products, which encompass many of the themes common to responsible AI frameworks. Geared toward an educator, family, and policy audience, the ratings are grounded in a set of eight core principles for how to think about AI and its impacts on adults and children:

- Put People First
- Promote Learning
- Prioritize Fairness
- Help People Connect
- Be Trustworthy
- Use Data Responsibly
- Keep Kids and Teens Safe
- Be Transparent and Accountable²¹

Similarly, Digital Promise, a global nonprofit working to expand educational opportunities through research and technology, lists four components in its AI evaluation framework: transparency, safety, ethics, and impact.²² And the National Institute of Standards and Technology includes these characteristics of trustworthy AI:

¹⁸ Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Buckingham Shum, S., Santos, O.C., Rodrigo, M.T., Cukurova, M., Bittencourt, I.I., & Koedinger, K.R. (2022). Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, 32:504–526. <https://doi.org/10.1007/s40593-021-00239-1>

¹⁹ Ibid.

²⁰ Ibid. Jurenka, I., Kunesch, M., McKee, K. R., Gillick, D., Zhu, S., Wiltberger, S., Phal, S. M., Hermann, K., Kasenberg, D., Bhoopchand, A., Anand, A., Pislár, M., Chan, S., Wang, L., She, J., Mahmoudieh, P., Rysbek, A., Ko, W.-J., Huber, A.,...& Ibrahim, L. (2024, May 14). *Towards responsible development of generative AI for education: An evaluation-driven approach*. Google DeepMind. <https://arxiv.org/abs/2407.12687>

²¹ Common Sense Media. (n.d.). *How we review and rate AI products*.

<https://www.common Sense Media.org/ai-ratings/how-we-rate-and-review>

²² Mills, K., Ruiz, P., & Lee, K. (2024, February 21). *Revealing an AI literacy framework for learners and educators*. Digital Promise.

<https://digitalpromise.org/2024/02/21/revealing-an-ai-literacy-framework-for-learners-and-educators/>

- Safe
- Secure and resilient
- Explainable and interpretable
- Privacy-enhanced
- Fair—with harmful bias managed
- Valid and reliable
- Accountable and transparent²³

Transparency and accountability are two of the toughest nuts to crack. AI algorithms are often difficult to understand, even for experts in the field, and large language models (LLMs) like ChatGPT are often referred to as "black boxes," or systems whose inner workings are invisible to the user.²⁴

Amara,²⁵ an AI ethics researcher, notes that a lot of the people who create these systems don't understand how the systems are working themselves.

As for accountability, Amara acknowledges the value of policy, but adds that AI systems are difficult to regulate even when rules have been created. She points to a New York City law in effect since 2023 that mandates an independent audit for any AI hiring tool. Companies must also disclose the use of any AI tools in the hiring process to prospective employees.²⁶ Amara questions the city's ability to monitor for violations and enforce any consequences. "We haven't figured out how to regulate these systems," she says.

Frameworks and principles are important not just for understanding, evaluating, and regulating products already on the market—they also need to be implemented at the design and development stages, before users are exposed to potential risks and harms. And it's crucial to be explicit about the values that inform the design of AI systems, platforms, and tools, because those same values inform the tools' outcomes. "If we build these tools with equity in mind, it will help us build equity in the rest of our systems. If we design them to track kids, they'll track kids," says Erin Mote, executive director of InnovateEDU. "Technology isn't neutral, and it's also built by humans."

²³ National Institute of Standards and Technology. (2023, January). *Artificial intelligence risk management framework*. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>

²⁴ Bagchi, S., & The Conversation U.S. (2023, May 26). Why we need to see inside AI's black box. *Scientific American*. <https://www.scientificamerican.com/article/why-we-need-to-see-inside-ais-black-box/>

²⁵ We have used a pseudonym at the researcher's request.

²⁶ New York City Department of Consumer Affairs. (n.d.). Automated employment decision tools. <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page>

Examples of AI Principles		
Common Sense Media	Digital Promise	NIST (National Institute of Standards and Technology)
Put People First	Transparency	Safe
Promote Learning	Safety	Secure and resilient
Prioritize Fairness	Ethics	Explainable and interpretable
Help People Connect	Impact	Privacy-enhanced
Be Trustworthy		Fair - with harmful bias managed
Protect Our Privacy		Valid and reliable
Keep Kids and Teens Safe		Accountable and transparent
Be Transparent and Accountable		

Cautionary notes

Before we turn to the meat of this white paper, there are a few cautionary notes to keep in mind:

Don't believe (all) the hype

Yes, generative AI represents a technological advancement, one that we couldn't have imagined just a generation ago. But there's a reason that the Substack blog called "AI Snake Oil," written by Princeton professor Arvind Narayanan and graduate student Sayash Kapoor, is popular.²⁷ A lot of what's being sold as AI isn't nearly as good as it's made out to be. Marketing speak has been around forever, and sales teams are out to make a profit. Caveat emptor.

Don't forget history

The tech industry is future-facing by design. That can be good for innovation, but bad for historical context. Although generative AI is new, algorithms have been in operation for decades, with many of the same bias problems that are popping up today.²⁸ And in education, technology has been seeping into the classroom for 100 years. Audrey Watters' book *Teaching Machines* traces the history of personalized learning back to Sidney Pressey's automated grading machines of the 1920s, followed by the more well-known inventions from B.F. Skinner. "Teaching machines are often decried these days as outdated behaviorist technologies, a disparagement that tends to overlook how much of Skinner's ideas—'conditioning,' in his terms, or 'nudging,' in more recent Silicon Valley parlance—have made their way into the classroom via our contemporary computing devices," she writes.²⁹

Don't confuse generative AI with actual intelligence

Generative AI is still artificial. Generative AI systems can't think like humans can. They can't reason, use judgment, or tell right from wrong. These models use predictive technology to produce outputs. The U.S. Department of Education's AI report explains it like this: "At its heart, AI is a highly advanced mathematical toolkit for building and using models. Indeed, in well-known chatbots, complex essays are written one word at a time. The underlying AI model predicts which next words would likely follow the text written so far; [generative] AI chatbots use a very large statistical model to add one likely word at a time, thereby writing surprisingly coherent essays."³⁰ More succinctly, Ashwin Sridhar, chief product and technology officer at Common Sense Media, describes generative AI as "auto-complete on steroids."

But most people don't know that's how the technology works. Amanda Bickerstaff, co-founder and CEO of AI for Education, which offers workshops, tools, and consulting on AI for educators and school systems, says there are significant misconceptions about this technology. "We ask, 'Is AI thinking?' And a lot of

²⁷ Narayanan, A., & Kapoor, S. AI snake oil. <https://www.aisnakeoil.com/>

²⁸ Turner Lee, N., Resnick, P., & Barton, G. (2019, May 22). *Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms*. Brookings Institution.

<https://www.brookings.edu/articles/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/>

²⁹ Watters, A. (2023.) *Teaching machines*. MIT Press.

³⁰ U.S. Department of Education, Office of Educational Technology. (2023, May). *Artificial intelligence and the future of teaching and learning: Insights and recommendations*. <https://www2.ed.gov/documents/ai-report/ai-report.pdf>

people, sometimes up to 50%, will say yes. And that's because of the way the technology has been released."

In education, this misconception is especially salient. Both teachers and students risk conflating the appearance of thinking with actual comprehension of subject matter.³¹

Don't trust generative AI as a source of truth

Because generative AI systems are built on predictive technology to return statistically relevant words, they can't be relied on to provide trusted, factual information. It's tempting to think of a chatbot as an authority; the design of these tools makes it seem like one. But any information a chatbot or other generative AI tool spits out must always be double- or even triple-checked against reliable sources of information. Generative AI, despite the nomenclature of "intelligence," is best used for creative exploration, not fact-finding and research.

Don't assume bigger is better

One of the biggest takeaways from Common Sense Media's first set of ratings and reviews: "The more data that an AI tool scrapes from the internet, the riskier it can be for users. In fact, the most successful AI tools that we reviewed are powered by limited, thoughtfully curated data sets and are designed for specific audiences or contexts."³²

Girard Kelly, senior counsel and head of privacy at Common Sense, comes at this from both an efficacy and a privacy perspective. One of his biggest concerns is the assumption that the more data a product takes in, the more effective it will be. "The truth of that remains to be seen," he says. And to protect users' privacy, he pushes back against the wide-net approach of the generative AI models on the market today. "If you don't need the data, don't collect it." Both data that a product is trained on, and data that a product collects from users to further refine its personalized answers or output, can impact privacy. As Kelly says, "We need to denormalize data collection by default."³³

Kelly's advice echoes one of the earliest—and most influential—papers on the perils of large language models, "On the Dangers of Stochastic Parrots." In it, the authors recommend mitigating against the risks of overly hegemonic viewpoints, encoded biases, and the difficulty of understanding just what is in the training data by "only creating datasets as large as can be sufficiently documented."³⁴

³¹ Blodgett, S.L., & Madaio, M. (2021, October 18). *Risks of AI foundation models in education*. Center for Research on Foundation Models, Stanford University. <https://crfm.stanford.edu/commentary/2021/10/18/blodgett-madaio.html>. For more information on additional risks posed by predictive technology, which is beyond the scope of our paper, see:

Akgun, S., & Greenhow, C. (2022). Artificial intelligence in education: Addressing ethical challenges in K–12 settings. *AI and Ethics*, 2:431–440. <https://doi.org/10.1007/s43681-021-00096-7>

³² Pizzo Frey, T. (2023, November 15.) *Championing responsible AI for kids and families*. Common Sense Media.

<https://www.common Sense Media.org/kids-action/articles/championing-responsible-ai-for-kids-and-families>

³³ King, J., & Meinhardt, C. (2024, February 22). *Rethinking privacy in the AI era: Policy provocations for a data-centric world*. Stanford University Institute for Human-Centered Artificial Intelligence (HAI).

<https://hai.stanford.edu/white-paper-rethinking-privacy-ai-era-policy-provocations-data-centric-world>.

³⁴ Bender, E.M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021.) On the dangers of stochastic parrots: Can language models be too big? *FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623.

<https://dl.acm.org/doi/pdf/10.1145/3442188.3445922>

How Americans view AI

As part of documenting the specific landscape of AI and education, we wanted to back up a bit, in order to understand first what Americans think about AI in general. Data from recent surveys on how people feel about AI in their lives mirrors what we heard in our conversations with educators, students, parents, and researchers, and is also similar to the wide range of both attitudes and knowledge among teachers in particular, as we'll discuss in the next section.


General attitudes

As we note in the introduction, viewpoints on AI range from *rah-rah* to *no thank you*. Recent public opinion surveys have reported substantial—and growing—concern among Americans, with some pockets of positivity. According to an August 2023 Pew Research Center survey, 52% of U.S. adults are more concerned than excited about AI in daily life—up from 38% in 2022. Just 10% say they are more excited than concerned; 36% feel a mix of excitement and concern.³⁵

A national opinion poll from Elon University, conducted in October 2023, also found a mix of optimism and wariness.³⁶ Only 17% of respondents said they feel more positive than negative about the impact of increased use of AI on people's daily lives. Respondents with a college degree and White respondents were more likely to say the impact of AI will be more negative.

Notably, more people in the Elon University survey believe AI will have a negative rather than positive impact on both K–12 and postsecondary education.

Overall, how will the increased use of AI affect the quality of people's daily lives?



Note: Those who did not answer are not shown

Source: Elon University poll of U.S. adults, Oct. 20–22, 2023

³⁵ Faverio, M., & Tyson, A. (2023, November 21). *What the data says about Americans' views of artificial intelligence*. Pew Research Center. <https://www.pewresearch.org/short-reads/2023/11/21/what-the-data-says-about-americans-views-of-artificial-intelligence/>

³⁶ Imagining the Digital Future Center. (2024, February 29). *The impact of artificial intelligence by 2024: National public opinion poll findings*. Elon University. <https://imaginingthedigitalfuture.org/wp-content/uploads/2024/02/AI2040-Report-public-opinion-poll-white-paper-1.pdf>

Awareness and knowledge

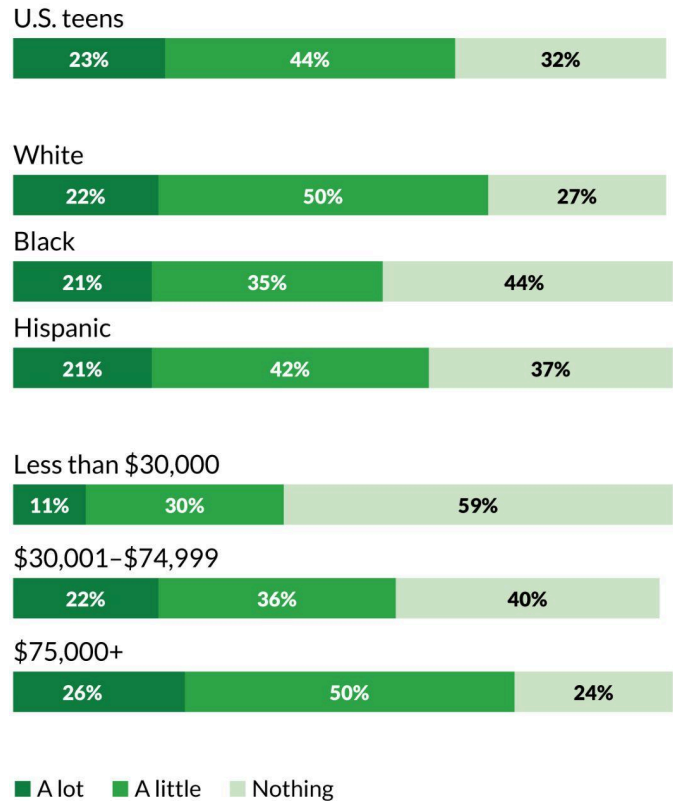
Americans have opinions about AI, but how much do they know about it? An August 2023 Pew Research survey found that the vast majority of Americans (90%) said they had heard at least "a little" about artificial intelligence, but only one in three said they'd heard "a lot" about it.³⁷ And while most people are aware of AI, fewer can identify AI in practice. Only 30% of U.S. adults correctly recognized all six examples of AI in everyday life that an earlier Pew survey asked about.³⁸

Among young people specifically, 23% of teens age 13 to 17 said they'd heard a lot about ChatGPT in particular, and 44% reported hearing a little about it. Awareness varies by race and ethnicity, as well as by household income, with more White teens and teens living in households making \$75,000 or more per year saying they've heard at least a little about ChatGPT, according to a Pew Research survey from September and October 2023.³⁹

Among parents, only 16% felt as though they had a detailed understanding of AI, according to a 2023 survey from the National Parents Union. Nearly two-thirds (62%) said they've heard little or nothing about AI tools in education, and fewer than half (49%) felt their child's school is well prepared to use AI tools. Yet more than two-thirds of parents believe that the potential benefits of using AI in K-12 education either outweigh or are equal to the potential downsides.⁴⁰

Most teens have heard of ChatGPT, but awareness varies by race and ethnicity, household income

% of U.S. teens ages 13 to 17 who say they have heard ___ about ChatGPT, an artificial intelligence (AI) program used to create text



Note: White and Black teens include those who report being only one race and are not Hispanic. Hispanic teens are of any race. There were not enough Asian American teens in the sample to analyze responses separately. Those who did not give an answer are not shown.

Source: Pew Research Center Survey of U.S. teens age 13-17 conducted Sept. 26-Oct. 23, 2023.

³⁷ Faverio, M., & Tyson, A. (2023, November 21). *What the data says about Americans' views of artificial intelligence*. Pew Research Center. <https://www.pewresearch.org/short-reads/2023/11/21/what-the-data-says-about-americans-views-of-artificial-intelligence/>

³⁸ Ibid.

³⁹ Sidoti, O., & Gottfried, J. (2023, November 16). *About 1 in 5 U.S. teens who've heard of ChatGPT have used it for schoolwork*. Pew Research Center.

<https://www.pewresearch.org/short-reads/2023/11/16/about-1-in-5-us-teens-whove-heard-of-chatgpt-have-used-it-for-schoolwork/>

⁴⁰ National Parents Union. (2023, October 17). *New poll: Parents see potential of artificial intelligence to enhance learning, but uncertainty persists*.

<https://nationalparentsunion.org/2023/10/17/new-poll-parents-see-potential-of-artificial-intelligence-to-enhance-learning-but-uncertainty-persists/>

A May 2024 Impact Research survey, commissioned by the Walton Family Foundation, reported similar numbers, with 69% of K–12 parents saying they think AI chatbots are a valuable tool for learning, and 57% saying chatbots have had a positive impact on their children. Parents of color in particular want to see AI chatbots used more in schools—64% of Black parents and 65% of Hispanic parents, vs. 55% of White parents.⁴¹

In one of our parent conversations, a New Jersey mother of a 10th grader who works in nonprofit finance has attended webinars about using AI to stay relevant in her career. She has found useful information but has also discovered a steep learning curve. "Honestly, the fact that we are using AI in workplaces, I think in schools there has to be a way to incorporate AI while making sure students are comprehending what they need to, to be able to think and analyze and build the skills they need to be successful," she says.

Among the possible benefits, according to the National Parents Union survey, parents rank AI-enhanced tutoring as having the most potentially positive effect, followed by initial feedback on homework from AI tools before students turn it in. In terms of teacher usage, parents see potential in AI-generated worksheets, AI-customized lesson plans based on student needs, and AI tools that analyze performance data and identify strengths and weaknesses.⁴²

A mother of a middle schooler in Brooklyn, New York, is on the wary end of the spectrum, extremely skeptical of generative AI usage for schoolwork. But she concedes possible benefits for teachers, such as drafting a lesson plan or assignment. "We all know how hard it is to start from a blank page," she says.

Usage

ChatGPT is the generative AI tool that Americans are most likely to have experimented with. A Consumer Reports survey that collected data in August and November 2023 found far more people trying ChatGPT than competitors from Google or Bing. Still, usage isn't widespread—69% of U.S. adult respondents reported not using *any* AI chatbot in the previous three months. The CR survey found that the most common use cases for chatbots were related to gathering and understanding information: to answer a question, to explain something, or to write, rewrite, or edit something.⁴³

Similarly, a February 2024 Pew Research survey found that most Americans still haven't used ChatGPT, despite a significant uptick since 2023. Usage varies, with adults age 18 to 29 and highly educated adults most likely to have used the chatbot.⁴⁴

⁴¹ Impact Research. (2024, May). *AI chatbots in schools: Findings from a poll of K–12 teachers, students, parents, and college undergraduates*. <https://www.waltonfamilyfoundation.org/learning/the-value-of-ai-in-todays-classrooms>

⁴² Ibid.

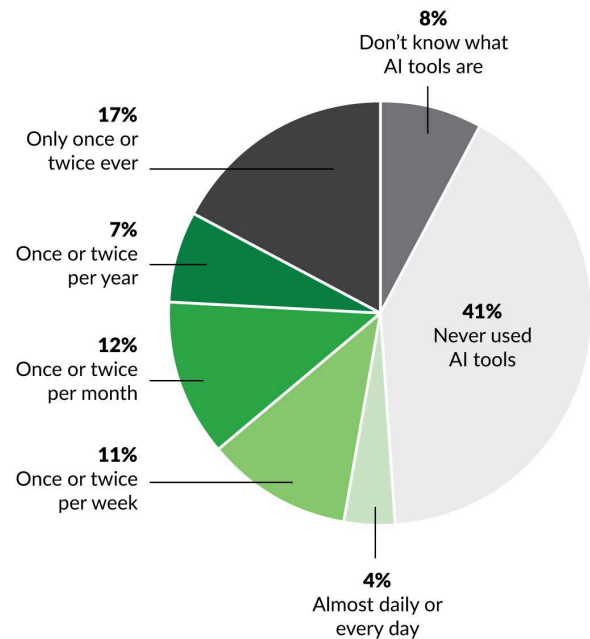
⁴³ Consumer Reports. (2024, January 30). *Consumer Reports surveys American consumers about how they use and think about text-based generative AI chatbots*. https://advocacy.consumerreports.org/press_release/consumer-reports-surveys-american-consumers-about-how-they-use-and-think-about-text-based-generative-ai-chatbots/

⁴⁴ McClain, C. (2024, March 26). *Americans' use of ChatGPT is ticking up, but few trust its election information*. Pew Research Center. <https://www.pewresearch.org/short-reads/2024/03/26/americans-use-of-chatgpt-is-ticking-up-but-few-trust-its-election-information/>

Use of generative AI tools increases further among younger age groups. A November 2023 survey from Common Sense Media⁴⁵ found that half of young people age 14 to 22 had used generative AI at some point in their lives. However, only 4% reported being daily users. The percentage of those who had ever used generative AI is similar across race/ethnicity, LGBTQ+ identity, gender, and age, but Black and Latino young people reported weekly or more generative AI use at approximately double the rate of White young people (22% and 18% vs. 10%, respectively).

The most commonly reported uses of generative AI were for getting information (53%) and brainstorming (51%). Several variations arise by demographic group, particularly in regard to race/ethnicity. Black young people are significantly more likely than their White peers to turn to generative AI to get information (72% vs. 41%), brainstorm ideas (68% vs. 42%), get help with schoolwork (62% vs. 40%), make pictures or images (51% vs. 24%), make sounds or music (42% vs. 7%), get help in their job (33% vs. 10%), and to write code (31% vs. 8%). Latino young people who have ever used generative AI are more likely than White young people to make pictures or images (39% vs. 24%), make sounds or music (27% vs. 7%), get help with their job (24% vs. 10%), get information (18% vs. 8%), and write code (18% vs. 8%).

Frequency of generative AI use among young people ages 14-22



A Pew survey done in fall 2023 found that almost one in five U.S. teens had used ChatGPT for schoolwork, and most (69%) who had heard of the tool said it's acceptable to use it to research new things. But only 39% said it's OK to use for solving math problems, and just 20% said it's OK for writing essays.⁴⁶ The students we spoke with (see our [takeaways below](#)) reported that some English and history teachers had mandated handwritten, paper-based essays to reduce the possibility that kids are using AI. (Interestingly, 54% of U.S. adults believe that generative AI programs should credit the source material they review, according to a February 2024 Pew survey.)⁴⁷

⁴⁵ Common Sense Media, Hopelab, and the Center for Digital Thriving at Harvard Graduate School of Education. (2024, June 3). *Teen and young adult perspectives on generative AI: Patterns of use, excitements, and concerns*.

<https://www.common Sense Media.org/sites/default/files/research/report/teen-and-young-adult-perspectives-on-generative-ai.pdf>

⁴⁶ Sidoti, O., & Gottfried, J. (2023, November 16). *About 1 in 5 U.S. teens who've heard of ChatGPT have used it for schoolwork*. Pew Research Center.

<https://www.pewresearch.org/short-reads/2023/11/16/about-1-in-5-us-teens-who've-heard-of-chatgpt-have-used-it-for-schoolwork/>

⁴⁷ Tyson, A., & Kennedy, B. (2024, March 26). *Many Americans think generative AI programs should credit the sources they rely on*. Pew Research Center.

<https://www.pewresearch.org/short-reads/2024/03/26/many-americans-think-generative-ai-programs-should-credit-the-sources-they-rely-on>

AI literacy among teachers

As we've seen, broad surveys have shown that most Americans still lack deep understanding of and knowledge about generative AI. That holds true in the domain of education as well, though usage has increased.

Before generative AI hit the scene, AI literacy in the education sphere—algorithms, machine learning, etc.—was often associated with teaching kids to code and preparing them for jobs in the tech industry. But as these technologies become part of everyday life, *everyone* needs to be AI literate, not just teachers and students in computer science and other technical classes. The concept of digital citizenship, which includes elements such as news literacy, media well-being, digital identity, and cyberbullying,⁴⁸ now must also thread AI literacy throughout all its components. AI's growing tentacles in American life have implications for the health of democracy, too—understanding how the technology works, and how it can be used ethically and unethically to shape and create content and sentiment, is essential to creating an informed populace with the agency to make decisions that impact their own lives as people and citizens.

When it comes to the need for AI literacy, educators, both classroom teachers and administrators, are the starting point. The rush to hype the benefits and opportunities of AI—*tailored curriculum in seconds! a tutor for every child!*—skips this crucial first step of learning the fundamentals.

If educators don't know the basics of the technology and how it can impact learning, they can't teach kids about it or make informed decisions about which educational tools are safe and effective.

If educators don't know the basics of the technology and how it can impact learning, they can't teach kids about it or make informed decisions about which educational tools are safe and effective.

"What schools know they need and have immediate demand for is teacher professional learning," says Alex Kotran, founder and CEO of the AI Education Project (also known as aiEDU), an AI literacy nonprofit. There's a common thread throughout the districts and schools that Kotran has worked with: If teachers and staff aren't informed about AI, none of the other initiatives that schools might be developing, from general AI policies to procurement guidelines, will work. In an aiEDU survey of educators, published in March 2024, an overwhelming majority said that professional development should "probably" or "definitely" include curriculum specifically designed to help them learn about the implications of AI.⁴⁹

Much like the concept of responsible AI, there's no commonly accepted standard for what it means to be AI literate, but consensus is emerging around certain elements. According to frameworks in development, the components of AI literacy include a basic understanding of how AI works, knowledge of how to use AI, and the

⁴⁸ Common Sense Education. (n.d.). *Everything you need to know about digital citizenship*. Common Sense Media. <https://www.commonsense.org/education/digital-citizenship>

⁴⁹ aiEDU. (n.d.) *Pulse survey report*. <https://www.aiedu.org/pulse>

capacity to critically evaluate AI tools and systems, including reflecting on risks and challenges.⁵⁰ We would add a fourth critical component for educators: understanding if and how AI tools and technology fit into the organizational structures, learning processes, and social realities of schools.

AI literacy components for educators

- A basic understanding of how AI works
- Knowledge of how to use AI
- Capacity to critically evaluate AI tools and systems
- Understanding how AI tools and technology fit into school structures and realities

For the most part, schools are still in the initial stages of the first and second categories. "We're still at early adoption," says Amanda Bickerstaff of AI for Education. "We just had 400 people join us live on a webinar, and only about 30% used AI regularly at all. And these are a self-selecting group of people coming to a webinar. There's a wide variance of adoption, of understanding, of familiarity, and of perceptions of [AI's] worth."

Recent surveys bear out Bickerstaff's observations, though familiarity is increasing. In the RAND Corporation's fall 2023 Teacher Omnibus, 44% of K-12 educators who took the survey had heard of AI tools but had never used them. Another 9% said this was the first time they were hearing about AI tools.⁵¹ Similarly, the aiEDU educator survey found that nearly half of respondents were "not at all familiar" or "slightly" familiar with generative AI. That percentage changes depending on grade level, with middle and high school teachers more likely to be "moderately familiar" or "extremely familiar" with the technology than those working at primary or elementary schools.⁵² Those numbers are starting to inch up, however. As of May 2024, 79% of K-12 teachers reported being at least somewhat familiar with ChatGPT in particular. Usage is likewise ticking up, with a substantial increase in K-12 teachers using ChatGPT at least once a week since February 2023, from 40% to 49%.⁵³

A high school social studies teacher in Huntington, New York, attended his first professional development session in early 2024. "It was very introductory, very eye-opening as well. I actually hadn't dabbled with ChatGPT until that point. It was eye-opening to see how powerful it can be," he says. A fourth grade math teacher in southern California hasn't played around with the technology much outside of her school's pilot use of an AI tutoring tool. "I know other people have, there's some people I know that live on ChatGPT, but I'm good with writing my own essays. So I don't use a lot of AI," she says. [See the Appendix for the [full Q&A](#) with each teacher.]

⁵⁰ aiEDU. (n.d.). *About us*. <https://www.aiedu.org/about>; Mills, K., Ruiz, P., & Lee, K. (2024, February 21). *Revealing an AI literacy framework for learners and educators*. Digital Promise.

<https://digitalpromise.org/2024/02/21/revealing-an-ai-literacy-framework-for-learners-and-educators>

⁵¹ RAND American Educator Panels. (Administered 2023, fall). *American teacher panel, 2023 teacher omnibus: AI questions data file*. RAND Corporation. <https://bentobento.info/surveys/235>

⁵² aiEDU. (n.d.) *Pulse survey report*. <https://www.aiedu.org/pulse>

⁵³ Impact Research. (2024, May). *AI chatbots in schools: Findings from a poll of K-12 teachers, students, parents, and college undergraduates*. <https://www.waltonfamilyfoundation.org/learning/the-value-of-ai-in-todays-classrooms>

Understanding how AI works covers everything from definitions of key terms and the ABCs of algorithms, to the ethics of facial recognition and the biases that are baked into these technologies. And no one needs an advanced degree in machine learning to get a handle on the basics. Building knowledge, however, can't be limited to learning definitions and concepts. Educators' professional development must include practical usage of AI tools. "In our point of view, literacy has to be directly connected to using the tools. Without hands-on keyboard experiences, it is extremely difficult to be able to have them more deeply understand what this technology can do, and understand where it's safe," Bickerstaff says.

The danger comes in jumping into use without basic understanding of the technology, and in expecting teachers to manage their own learning. More than 60% of teachers report using ChatGPT in their work at least "once in a while."⁵⁴ And yet only 23% of district leaders in the fall 2023 American School District Panel, also from the RAND Corporation, reported that their district had provided training to teachers about the use of generative AI, such as ChatGPT.⁵⁵ In a May 2024 post, the Center for Reinventing Public Education suggested a national teacher-training effort akin to initiatives in Asia.⁵⁶

During this school year (2023–2024), how often do you typically use ChatGPT in your work as a teacher?

Almost every day

2%

A few times a week

6%

Once a week

15%

Every once in a while

42%

Never

34%

Source: RAND American Educator Panels, American School Leader Panel, 2023 ASDP Survey data file, RAND Corporation, Santa Monica, CA. March 2024


⁵⁴ RAND American Educator Panels. (Administered 2023, fall). *American teacher panel, 2023 teacher omnibus: AI questions data file*. RAND Corporation. <https://bentobento.info/surveys/235>

⁵⁵ RAND American Educator Panels. (Administered 2023, fall). *American school leader panel, 2023 ASDP survey data file*. RAND Corporation. <https://bentobento.info/surveys/247>

⁵⁶ Lake, R. (2024, May). *AI is coming to U.S. classrooms, but who will benefit?* The Center on Reinventing Public Education. <https://crpe.org/ai-is-coming-to-u-s-classrooms-but-who-will-benefit/>

"We are in the infancy of AI literacy, and educational leaders must be mindful not to become passengers in this effort. If not, things will be done to us," says Jean-Claude Brizard, president and CEO of Digital Promise. Brizard wants educational leaders, especially at the state level, to take a leadership role in order to inform the work of districts. "Local education associations right now seem to be on their own, and in some cases teachers are on their own trying to figure this out," he says. A related concern is that opportunities for educator training will no doubt be unequally distributed, which will have inequitable downstream effects on students.

Has your district provided training to your teachers about use of generative artificial intelligence (like ChatGPT?)



Source: RAND American Educator Panels, American School Leader Panel, 2023 ASDP Survey data file, RAND Corporation, Santa Monica, CA. March 2024

The state of school policy

Brizard's remark that we're still "in the infancy" of AI literacy can be applied to school AI policy as well.

In the RAND American School Leader Panel survey, almost half of respondents said their district had not adopted a policy specifically about students' use of generative AI. Nearly a third said they were working on one, and another 17% said their acceptable use policy already covers generative AI. Only 5% reported adopting a specific generative AI policy.⁵⁷ Similarly, 85% of respondents in RAND's Teacher Omnibus Survey said their district had no policy on AI usage, or they didn't know if their district had one.⁵⁸ Six months later, the numbers weren't that different, with 78% of teachers reporting that their K-12 school had no AI policy or that they didn't know if one existed.⁵⁹ A 2024 Center for Democracy and Technology survey of middle school and high school teachers found some positive momentum on school policy, but only insofar as decisions around general permissions or bans of ChatGPT and other generative AI tools.⁶⁰ Teachers we spoke with as part of our research for this paper said they didn't know of any specific AI guidelines from their schools. In short, we have a long way to go toward defining and adopting clear guidelines for generative AI use in schools.

Parents are also often left in the dark. One mother of a high school student in Maryland says there has been no communication directly from her child's school or district regarding generative AI policies. She says she assumes "they're either working on it, or they've given up and have decided it's inevitable and they can't do anything about it." Another parent in Brooklyn says that as far as she knows, her seventh grade daughter isn't using AI in school. "Though I am realizing I don't have a durable definition of AI in my mind," she adds. "To the extent that I have heard anything about AI in school, it's been a warning: No one should be submitting essays using ChatGPT. I don't know how they are policing." A 2023 Center for Democracy and Technology survey found that only 4 in 10 parents said they or their child had received guidance from their school on how to use generative AI responsibly for schoolwork and within school rules.⁶¹

Has your district adopted a policy specifically about students' use of generative artificial intelligence (like ChatGPT?)

No, we have not adopted a policy specifically about students' use of generative AI

47%

No, we're working on a policy specifically about students' use of generative AI, but have not yet adopted it

31%

No, a specific policy is not needed because our district's acceptable use policy covers students' use of generative AI

17%

Yes, we've adopted a policy specifically about students' use of generative AI

5%

Source: RAND American Educator Panels, American School Leader Panel, 2023 ASDP Survey data file, RAND Corporation, Santa Monica, CA, March 2024

⁵⁷ Ibid.

⁵⁸ RAND American Educator Panels. (Administered 2023, fall). *American teacher panel, 2023 teacher omnibus: AI questions data file*. RAND Corporation. <https://bentobento.info/surveys/235>

⁵⁹ Impact Research. (2024, May). *AI chatbots in schools: Findings from a poll of K-12 teachers, students, parents, and college undergraduates*. <https://www.waltonfamilyfoundation.org/learning/the-value-of-ai-in-todays-classrooms>

⁶⁰ Dwyer, M., & Laird, E. (2024, March). *Up in the air: Educators juggling the potential of generative AI with detection, discipline, and distrust*. Center for Democracy & Technology.

<https://cdt.org/wp-content/uploads/2024/03/2024-03-21-CDT-Civic-Tech-Generative-AI-Survey-Research-final.pdf>

⁶¹ Laird, E., Dwyer, M., & Grant-Chapman, H. (2023, September). *Off task: How machine learning algorithms in education are facilitating mass student surveillance*. Center for Democracy & Technology. <https://cdt.org/wp-content/uploads/2023/09/091923-CDT-Off-Task-web.pdf>

Digital Promise's Brizard identifies two ways that state education leaders in particular can be leaders on AI literacy: first by informing educators and working with them to determine how best to use AI tools, and second through the procurement process. He recalls a meeting with superintendents of large urban schools during which they asked him how they can get developers to adhere to certain standards or meet certain requirements in their products. Brizard turned the question back on the superintendents by inquiring whether they asked the developers directly for what they wanted. "People may have edtech strategies, and I'm using the term interchangeably because AI for me is edtech, but they don't have a mechanism for demanding certain levels of efficacy, of privacy, etc. So we are on a campaign right now to get folks to understand what they should look at in procurement, because that's the only way we are going to drive the market to do quickly what it needs to do and ought to be doing for our kids," Brizard says.

We are on a campaign right now to get folks to understand what they should look at in procurement, because that's the only way we are going to drive the market to do quickly what it needs to do and ought to be doing for our kids. — Jean-Claude Brizard

In Chile, we see an example of how the public sector can drive responsible AI. The World Privacy Forum's December 2023 report "Risky Analysis," which contains an international review of AI governance tools, explains how the country's "bidding and quality assurance requirements for government acquisition of AI systems stress the importance not only of evaluating the system's impacts on equity, but of evaluating the equity metrics themselves."⁶²

There are certainly bright spots. Chantell Manahan, director of technology at the Steuben County school district in Angola, Indiana, says her district is working with the teachers union on a set of AI beliefs. She has also been beating the procurement practices drum for years, well before generative AI hit the scene. Manahan's district provides teachers with workshop accounts to try new tools. If they want to request approval for official use, they fill out a form, and the technology department reviews and makes a recommendation. Some questions Manahan asks when reviewing include, "What's the privacy policy, the terms of use? Is this a tool for everyone? For teachers only? For 13 plus?" She knows that her work doesn't stop with approving or rejecting an AI tool, though. "Then we have to teach teachers how to effectively use those tools," she says.

A Digital Promise report published in February 2024 counted seven states that had issued guidance documents on AI in education: California, North Carolina, Ohio, Oregon, Virginia, Washington, and West Virginia. While that's a good start, and these states should be commended for taking the lead, the report noted that these documents generally describe the *need* for policies, rather than detailing the policies themselves. As the report concluded, "[T]here is clearly work to be done."⁶³ Since then, at least nine more states have published documents, but guidance remains generic. "It's important to build in a school- and district-centered context—this is why the

⁶² Kaye, K., & Dixon, P. (2023, December). *New report: Risky analysis: Assessing and improving AI governance tools*. World Privacy Forum. <https://www.worldprivacyforum.org/2023/12/new-report-risky-analysis-assessing-and-improving-ai-governance-tools>

⁶³ Roschelle, J., Fusco, J., & Ruiz, P. (2024, February 26). *Review of guidance from seven states on AI in education*. Digital Promise. <https://digitalpromise.dspacedirect.org/server/api/core/bitstreams/654a148a-25c4-4eec-bb8e-2fd1948130a5/content>

specificity isn't there at the state level. We've found that doing this work well requires significant time and resources," says Brizard.

Policy approaches in the U.S., across the Atlantic, and beyond

When OpenAI, for a brief moment in 2023, ousted its CEO and it looked like the company might implode, tech journalist Karen Hao had one big takeaway. "I think the biggest lesson is that self-regulation just doesn't work," she said in an interview with Apple News's Big Think.⁶⁴

Different countries are at different stages of legislation and regulatory action concerning technology, with most still in a nascent stage with respect to AI. We discuss some of the most significant developments below. The European Union, United Kingdom, and Australia offer some models for tech regulation, and the E.U. in particular has developed the first comprehensive framework for regulating digital products.

We also note that rather than developing new legislation specific to AI or generative AI, some governments are relying on existing laws to regulate and address core technology risks around privacy and consumer protection. (We see this in schools as well: In a study from Digital Promise, district leaders often preferred to integrate policy on AI into their broader policy documents.⁶⁵) And we acknowledge that AI regulation doesn't fit neatly into a one-size-fits-all framework—different countries have different views on AI, and this represents a small subsection of global thought.

United States

Much like school, district, and state-level policy on AI in education, U.S. federal AI policy overall is still in development. As of July 2024, Congress has not passed any AI-focused legislation.

There are four relevant policy developments to be aware of:

- **COPPA Law**

The Children's Online Privacy Protection Act of 1998 and its associated FTC rule, known as COPPA, is the 25-year-old main national law on the books that protects the privacy of children under 13. The Federal Trade Commission announced proposed modifications to its rule in early 2024, five years after the organization last requested comment and 11 years after the last round of updates.⁶⁶ Suggested modifications include additional layers of consent for using and sharing personal information, and a school authorization exception from consent requirements, in the form of a detailed written agreement between a technology provider and a school.⁶⁷

⁶⁴ Hao, K., & Chapman-Smith, R. (2023, November 27). *The chaos inside OpenAI*. Big Think.

<https://bigthink.com/videos/openai-chaos-explained/>

⁶⁵ Armstrong, A., Karim, S., & Ruiz, P. (2024, February 29). *How school districts are integrating generative AI into their policies*. Digital Promise.

<https://digitalpromise.org/2024/02/29/how-school-districts-are-integrating-generative-ai-into-their-policies/>

⁶⁶ Federal Trade Commission. (2024, January 11). *Children's online privacy protection rule*. Federal Register.

<https://www.federalregister.gov/documents/2024/01/11/2023-28569/childrens-online-privacy-protection-rule>

⁶⁷ Folks, A. (2024, February 16). *Still growing up: Top takeaways from the FTC's proposed COPPA rule update*. International Association of Privacy Professionals. <https://iapp.org/news/a/still-growing-up-top-takeaways-from-the-ftcs-proposed-coppa-rule-update/>; Common Sense Media. (2024). *COPPA 2.0 one-pager*.

<https://www.common sense media.org/sites/default/files/featured-content/files/coppa-2.0-one-pager-2024.pdf>

- **KOSPA Legislation**

In Congress, there are proposals to update the privacy rules and add new protections for youth safety. Introduced in 2024, the Kids Online Safety and Privacy Act (KOSPA) is the combination of two major pieces of online safety legislation for children, the Kids Online Safety Act (KOSA) and the Children and Teens' Online Privacy Protection Act (COPPA 2.0). The combined bill would require social media and gaming companies to design their platforms to prioritize children's and teens' safety and privacy, requiring them to remove design features that work to maximize kids' engagement. Companies would also be required to provide safeguards that are readily accessible and easy to use, including the strongest privacy and safety settings by default, parental tools to supervise use by children, and mechanisms to report harms. While the bill doesn't include considerations for generative AI, it would require platforms and applications to provide notification of the use of algorithms.⁶⁸

The bill, which passed the U.S. Senate on July 30, 2024, and will likely be taken up by the House in the fall, would also provide updates to COPPA. These updates include expanding protections to children under age 17 (up from age 13), banning targeted advertising to all covered minors, and minimizing the amount of data collected on minors.⁶⁹ The proposals do not directly address generative AI.

- **White House Executive Order**

In October 2023, the White House published an executive order (EO) on the development and use of AI—the only one of the four policy developments mentioned here that explicitly focuses on AI. While the EO doesn't create a new law, it seeks to use existing authority to create rules and procedures that control how the federal government and its entities approach AI. It includes everything from combating fraudulent AI-generated information and strengthening privacy guidance to addressing algorithmic discrimination and maximizing AI's benefits for workers. The order instructed the Department of Education to develop its own AI policies to "address safe, responsible, and nondiscriminatory uses of AI in education, including the impact AI systems have on vulnerable and underserved communities." The order also mandated an "AI toolkit" for education leaders to implement recommendations, including "appropriate human review of AI decisions, designing AI systems to enhance trust and safety and align with privacy-related laws and regulations in the educational context, and developing education-specific guardrails."⁷⁰ In July 2024, the department released its first report in response, which provides guidance for product developers. The recommendations focus on designing for teaching and learning, providing evidence of impact, advancing equity, ensuring safety and security, and promoting transparency.⁷¹

⁶⁸ Common Sense media. *Fact sheet—KOSPA*.

https://www.common sense media.org/sites/default/files/featured-content/files/kospa-one-pager-2024-updated-7.24.24_0.pdf

⁶⁹ Children and Teens' Online Privacy Protection Act, S.1418, 118th Cong. (2023). <https://www.congress.gov/bill/118th-congress/senate-bill/1418>; Children and Teens' Online Privacy Protection Act, H.R.7890, 118th Cong. (2023). <https://www.congress.gov/bill/118th-congress/house-bill/7890>

⁷⁰ The White House. (2023, October 30). *Executive order on the safe, secure, and trustworthy development and use of artificial intelligence*.

<https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>

⁷¹ U.S. Department of Education, Office of Educational Technology. (2024, July). *Designing for education with artificial intelligence: An essential guide for developers*. <https://tech.ed.gov/designing-for-education-with-artificial-intelligence/>

- **Digital Equity Act**

The Digital Equity Act, passed in 2021 as part of the Infrastructure Investment and Jobs Act, provides \$2.75 billion in digital equity capacity grants. These grants will offer schools and other anchor institutions funding to step up their AI digital literacy offerings for teachers, students, and caregivers to support the equitable use of AI in schools.⁷² In December 2023, Delaware Representative Lisa Blunt Rochester introduced an amendment to codify AI literacy as a component of digital literacy. The bill states that "effective AI initiatives encompass not only technical training but also include comprehensive education about the potential benefits and risks associated with AI technologies."⁷³

Amina Fazlullah, senior director of equity policy at Common Sense Media, says that tech policy in general, which should include AI-specific provisions, needs to focus on privacy, data security, platform responsibility, and AI literacy. Early AI-specific policies currently include two main components: auditing and enforcement. First, she says, "We need guardrails around the uses of AI to make sure those uses are not harmful. Companies build their algorithms to benefit their bottom line. Those goals don't necessarily match up to protect consumers—policy can step in to fix this mismatch," she says. Audits of AI systems and accompanying risk assessment frameworks, which the European Union is now in the process of designing, can help evaluate the potential for harm before products are released. "If we don't create rules of the road," says Fazlullah, "products will be released into the marketplace that could harm students and families."

Then we need mechanisms that would enforce those existing policies. Companies shouldn't avoid liability just by blaming an algorithm for making a decision, Fazlullah says. This notion of enforcement goes back to the accountability theme of responsible AI: Who is responsible, and how can government entities create feasible enforcement systems to hold those actors' feet to the fire?

Much like state education departments, some states have gone further than the federal government in terms of protecting privacy. California has a new agency handling data privacy, for example. In the 2023 legislative session, at least 25 states, Puerto Rico, and the District of Columbia introduced bills concerning AI, and 18 states and Puerto Rico adopted resolutions or enacted legislation.⁷⁴ Some include language on data privacy and accountability, and there are also related efforts to tackle other AI-connected issues, such as deepfakes and AI literacy for educators and students.

⁷² National Telecommunications and Information Administration. *Infrastructure investment and jobs act new funding initiatives*. <https://www.doi.gov/sites/doi.gov/files/4.e-ntia.pdf>

⁷³ Artificial Intelligence Literacy Act of 2023, H.R. 6791, 118th Cong. (2023). <https://bluntrochester.house.gov/uploadedfiles/20231215aibill.pdf>; <https://www.congress.gov/bill/118th-congress/house-bill/6791/all-info>

⁷⁴ Artificial Intelligence 2023 Legislation. (2024, January 12). National Conference of State Legislatures. <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2023-legislation>

European Union

The European Union's Artificial Intelligence Act (AIA) was formally approved in March 2024 and officially became law in August. Its various provisions will come into effect over the next two years.⁷⁵ The E.U.'s AIA sets a first-of-its kind global policy standard, while most other governments to date are either taking a voluntary approach to standard-setting or developing limited applications of AI regulation under existing laws.⁷⁶ The legislation identifies four different risk levels and sets requirements accordingly—the riskier the category, the more stringent the requirements.

Notably, while the AIA does not explicitly focus on children's rights or safety, specific uses of AI that could ultimately impact children, such as in education and toys, will meet the threshold for high risk. Other high-risk areas include health care, banking, critical infrastructure, and justice and democratic processes. High-risk systems must register with a centralized database, adhere to audit requirements, assess and reduce risks, maintain use logs, adhere to standards of transparency and accuracy, and ensure human oversight. Citizens have the right to submit complaints about AI systems and receive explanations about AI-based decisions that affect their rights.⁷⁷

The new E.U. act's rules also ban AI systems that pose an "unacceptable risk." These include untargeted scraping of facial images to create facial recognition databases, emotion recognition in the workplace and schools, and predictive policing when based solely on profiling a person.⁷⁸

The E.U. AI Act was in the works years before generative AI arrived on the scene. Following the advent of this new technology and the release of popular tools such as ChatGPT, the E.U. ultimately agreed to carve out distinct guardrails for so-called "general purpose" AI systems in order to balance the importance of innovation with the need to mitigate known—and unknown—risks. The threshold for regulating these systems, which include many generative AI tools trained on large language models, is high—it's estimated that only two or three existing general purpose AI systems currently on the market meet the E.U. AIA threshold for regulation (such as Google DeepMind's Gemini Ultra and OpenAI's GPT-4).⁷⁹ Further, the transparency requirements for these systems are minimal. Meanwhile, the U.S. executive order uses an even higher threshold, and no current generative AI models appear to meet its standard for regulation.

Before the AI Act began to take shape, the E.U.'s Digital Services Act (DSA), passed in 2022 and now in full effect, represented the world's first comprehensive online safety legislation. To that end, the DSA set a model for regulating the digital world, based on transparency and the capacity for risk. The DSA also has the potential to serve as an early use case for regulating generative AI, even before the AI Act comes into full effect, because some generative AI systems and products are likely to be covered by its obligations.⁸⁰

⁷⁵ Hickman, T., Lorenz, S., Teetzman, C., & Jha, A. (2024, July 16). *Long awaited EU AI Act becomes law after publication in the EU's Official Journal*. White & Case.

<https://www.whitecase.com/insight-alert/long-awaited-eu-ai-act-becomes-law-after-publication-eus-official-journal>

⁷⁶ European Parliament. (2024, March 8). *Artificial intelligence act: MEPs adopt landmark law*.

<https://www.europarl.europa.eu/news/en/press-room/20240308IPR19015/artificial-intelligence-act-meps-adopt-landmark-law>

⁷⁷ Jain, R. (2024, March 15). *The European Union's AI act: What you need to know*. Holland & Knight LLP.

<https://www.hklaw.com/en/insights/publications/2024/03/the-european-unions-ai-act-what-you-need-to-know>

⁷⁸ European Parliament. (2024, March 8). *Artificial intelligence act: MEPs adopt landmark law*.

<https://www.europarl.europa.eu/news/en/press-room/20240308IPR19015/artificial-intelligence-act-meps-adopt-landmark-law>

⁷⁹ Internal Common Sense Media analysis.

⁸⁰ Vermeulen, M., & Lemoine, L. (2024, February 12). *From ChatGPT to Google's Gemini: When would generative AI products fall within the scope of the Digital Services Act?* Media@LSE, LSE Department of Media and Communications.

Recently, political momentum in the E.U. has coalesced around combating the harms to children of "addictive design"—features of applications or devices that are intended to hook users, and particularly minors, into frequent use.⁸¹ These AI-driven features could continue to be a significant focus of DSA enforcement, or even a topic of potential legislation under a new E.U. administration. For example, when TikTok announced a new feature in France and Spain that would reward users for performing certain tasks on the app, the European Commission opened an investigation into whether TikTok had conducted an advance risk assessment, which is required by the DSA, to gauge the potential systemic risks of the functionalities for minors' mental health.⁸² TikTok immediately suspended the feature and paused the rollout in additional E.U. countries.⁸³

Legislation, of course, is only as good as its implementation, and one critical area to watch, especially as newly elected commissioners take office this fall after June elections, will be how the European Commission and different European countries set up and staff regulatory bodies that enforce the AIA's rules. European regulators will need to keep track of existing AI models while keeping abreast of new developments. What sets the E.U. AI Act apart is that compliance is mandatory, and comes with significant fines for violation, says Sara Egozi, head of global strategy and policy at Common Sense.

United Kingdom and Australia

Unlike the E.U., the U.K. and Australia have not yet developed AI-specific legislation to date. Rather, they've taken limited steps that include assessing potential regulation of generative AI under existing online safety laws. (In both countries, the law is called the Online Safety Act.)

The U.K. has taken a fairly pro-innovation approach to date, employing a potential cross-sector framework for regulating AI that involves several existing regulatory entities, along with voluntary safety and transparency measures. Nearly 350 staff have been dedicated thus far to enforcement of the U.K.'s Online Safety Act, with another 100 expected to be hired this year,⁸⁴ and the country's telecommunications regulator, Ofcom, is holding a series of consultations that will inform the eventual rules, which are expected in 2025.⁸⁵ In mid-April 2024, however, the U.K. reportedly started drafting legislation to govern AI,⁸⁶ and King Charles III announced potential upcoming regulations in a July speech.⁸⁷ Under the new Labour government, a comprehensive AI bill is expected sometime in 2025.

<https://blogs.lse.ac.uk/mediase/2024/02/12/from-chatgpt-to-googles-gemini-when-would-generative-ai-products-fall-within-the-scope-of-the-digital-services-act/>

⁸¹ Common Sense Media. (n.d.) *Can media be addictive?*

<https://www.commonsense.org/education/digital-citizenship/lesson/can-media-be-addictive>

⁸² European Commission. (2024, April 22). *Commission opens proceedings against TikTok under the DSA regarding the launch of TikTok Lite in France and Spain, and communicates its intention to suspend the reward programme in the EU.*

https://ec.europa.eu/commission/presscorner/detail/en/ip_24_2227

⁸³ O'Carroll, L. (2024, 24 April). TikTok reward-to-watch feature suspended after EU threats to block it. *The Guardian*.

<https://www.theguardian.com/technology/2024/apr/24/tiktok-reward-to-watch-feature-suspended-after-eu-threats-to-block-it>

⁸⁴ Criddle, C. (2024, January 8). Ofcom poaches Big Tech staff in push to enforce new internet curbs. *Financial Times*.

<https://www.ft.com/content/19e9da57-5da4-40bc-989d-9820bf3d2aff>

⁸⁵ Office of Communications (Ofcom). (2023). *Protecting people from illegal content online.*

<https://www.ofcom.org.uk/consultations-and-statements/category-1/protecting-people-from-illegal-content-online>

⁸⁶ Milligan, E. (2024, April 15). *UK starts drafting AI regulations for most powerful models.* Bloomberg.

<https://www.bloomberg.com/news/articles/2024-04-15/uk-starts-drafting-ai-regulations-for-most-powerful-models>

⁸⁷ Prime Minister's Office, 10 Downing Street, and His Majesty King Charles III. (2024, July 17). *The king's speech 2024.*

<https://www.gov.uk/government/speeches/the-kings-speech-2024>

In January 2024, Australia announced its intention to implement mandatory safeguards for high-risk AI use cases. Requirements, developed through consultation with industry and the public, are expected to revolve around testing and audits, transparency, and accountability.⁸⁸

Governments around the world will likely continue to grapple with the fundamental tension between addressing risk and fostering innovation, says Egozi. "But innovation and risk mitigation are not mutually exclusive. Rather, when it comes to kids, families and consumers stand to benefit most from responsibly built AI products, whether intended for education in the classroom or use at home."

Governance models

Once guidelines and enforcement mechanisms have been established, the work isn't necessarily done. The implementation layer can be just as challenging. The World Privacy Forum's international review found that AI governance tools, or tools and techniques to evaluate AI systems for inclusiveness, fairness, explainability, privacy, and other responsible AI elements, "too often lack meaningful oversight and quality assessments." The review found that 38% of AI governance tools mention, recommend, or incorporate one of three measures that have been shown to be problematic. The report also found that standards and guidance for assessment aren't consistent,⁸⁹ which makes sense given the early stage of this ecosystem's development.

Pizzo Frey also points out that audit requirements, which frequently show up as part of governance frameworks, often lack specificity. "One of the biggest areas that needs to be addressed is precise scoping and documentation requirements," she says. This is especially important because tech companies tend to advocate for audits to be executed internally, rather than by third parties, which can present conflicts of interest. A recent paper warned of the "dangers of relying on a magic incantation of 'audits' to address algorithmic accountability challenges" without providing for the effective participation of third parties. "A future where AI audits are conducted by company employees and consulting companies alone fails to learn the lessons of how third parties have generated robust oversight of AI systems. Such systems are diverse, impacting a range of industries and stakeholders, and it will require more third party participants, representing a variety of communities and making use of a wide set of tools, to concretely assess a wider range of the harms involved for the implicated stakeholders," the authors wrote.⁹⁰

⁸⁸ Boyle, N., & Horder, A. (2024, January 31). *AI regulation in Australia: What we know and what we don't*. DLA Piper.

<https://www.dlapiper.com/en/insights/publications/2024/01/ai-regulation-in-australia-what-we-know-and-what-we-dont>

⁸⁹ Kaye, K., & Dixon, P. (2023, December). *New report: Risky analysis, assessing and improving AI governance tools*. World Privacy Forum.

<https://www.worldprivacyforum.org/2023/12/new-report-risky-analysis-assessing-and-improving-ai-governance-tools>

⁹⁰ Raji, I.D., Xu, P., Honigsberg, C., & Ho, D. (2022, July). Outsider oversight: Designing a third-party audit ecosystem for AI governance. *AIES '22: Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, 557–571. <https://doi.org/10.1145/3514094.3534181>

Opportunities

We approach generative AI's potential for positive impact in education with caution. This technology is still experimental, and we have yet to understand its long-term effects. AI also carries risks and challenges, as we detail in the next section. Still, there are promising opportunities, especially if designed from the beginning with equity in mind, and developed with care and intention.

We begin from the premise that AI technology in education should never diminish students or educators as humans.⁹¹ In fact, strong human systems are required for ethical, effective, and safe AI deployment. As Common Sense highlighted in its AI principles and initial product reviews, AI is sociotechnical, which means that the technology cannot be separated from the humans and human-created processes that inform, shape, and develop its use.⁹²

We have divided opportunities into three main types: student-supporting, teacher-supporting, and system-supporting. We have subdivided further within each type. These opportunities are not intended to be an exhaustive or dispositive list; rather, they represent themes that emerged from our interviews for this paper. We chose to illustrate them by sharing the experiences and perspectives of those working in the field, whether in the classroom, the lab, or the product suite. That choice is yet another way to elevate the importance of human involvement.

Student-supporting

We highlight three opportunity areas with student-supporting AI tools. The first, adaptive learning, is what many people think of when they hear "edtech." The idea of a bespoke tutor perfectly matched to each individual child is often presented as the holy grail of digital-enhanced learning. We're not there yet, but we recognize progress in tutoring systems.

The next two areas we focus on, creativity and project-based learning, fit with one of the conclusions from Common Sense's first product ratings and reviews. "Generative AI is best for fiction, not fact. Consumers—and kids especially—must understand that generative AI tools are best used for creative exploration and are not designed to consistently give factual, truthful responses to questions."⁹³ The examples we highlight represent creative applications of generative AI to build cognitive and collaborative skills.

⁹¹ Aiken, R. M., & Epstein, R.G. (2000). Ethical guidelines for AI in education: Starting a conversation. *International Journal of Artificial Intelligence in Education*, 11, 163–176.

https://www.researchgate.net/profile/Richard-Epstein-8/publication/228600407_Ethical_guidelines_for_AI_in_education_Starting_a_conversation/links/5523ef290cf2c815e073e5b0/Ethical-guidelines-for-AI-in-education-Starting-a-conversation.pdf

⁹² Common Sense Media. (n.d.). *How we review and rate AI products*.

<https://www.common sensemedia.org/ai-ratings/how-we-rate-and-review>

⁹³ Pizzo Frey, T. (2023, November 15.) *Championing responsible AI for kids and families*. Common Sense Media.

<https://www.common sensemedia.org/kids-action/articles/championing-responsible-ai-for-kids-and-families>

Adaptive learning

For many years, well before the emergence of generative AI, adaptivity has been touted as one way technology can improve learning.⁹⁴ Research has shown that students can learn more with strategies that adapt to their needs, as opposed to a fixed sequence of examples and problems.⁹⁵ Generative AI may further advance learning improvements by meeting students where they are and morphing to fit their needs.⁹⁶ Interestingly, two of the three tools we highlight here have chosen *not* to use generative AI directly with students for adaptive learning, citing concerns about safety, accuracy, and efficacy. Instead, they use more restricted forms of AI.

Eedi is a math tutoring system that claims to close knowledge gaps by identifying misconceptions and reteaching. As the founders worked on the tool in 2020, they looked to customer service for inspiration. One-to-one tutoring was the gold standard, says Eedi CEO Ben Caulfield, and they wanted to find a way to reap those benefits using fewer human resources. So they looked at service platforms that started out with automated answers and pivoted to human agents if customers couldn't resolve their issues.

With Eedi, diagnostic questions written by human experts assess students' knowledge and identify gaps, and a chat system helps them learn. Video lessons provide worked examples by human tutors. If a student hits a roadblock, a qualified human teacher steps into the chat live (a fee-based feature available with Eedi Plus). Once the student has made progress, the tutor directs them back to the lesson.

While Eedi uses machine learning models to predict how students will answer questions and to choose the next best question for them, the company has held off on using generative AI directly with students. Eedi uses large language models (LLMs) to suggest responses to human tutors, but the tutors have complete control over whether to use the AI-generated response, edit it, or write something entirely different. Retaining human oversight guards against the risk of inaccuracies and hallucinations, and provides built-in moderation. "Even 99% accurate may not be enough, and red-teaming [testing a system for flaws and cracks] is never enough," says Caulfield.

Eedi makes its data sets, which include 120 million student answers, available to researchers once a data-sharing agreement has been signed. The data is sent through a secure, time-limited link and contains no personal information; the only thing Eedi can reasonably guess at is an indication of age, based on the lessons. For their research studies, they collect demographic information via an explicit request for permission. Caulfield says that research, which has so far shown positive results in small sample sizes⁹⁷ (a study is underway with a larger cohort), is key to their mission. Without naming names, he critiques "solutions with little evidence of the real effect on student learning."

⁹⁴ Aleven, V., McLaughlin, E. A., Glenn, R. A., & Koedinger, K. R. (2016). Instruction based on adaptive learning technologies. In Mayer, R.E. & Alexander, P.A., *Handbook of research on learning and instruction*, 522–560. ISBN: 113883176X as cited in U.S. Department of Education, Office of Educational Technology. (2023, May). *Artificial intelligence and the future of teaching and learning: Insights and recommendations*. <https://www2.ed.gov/documents/ai-report/ai-report.pdf>

⁹⁵ Najjar, A.S., Mitrovic, A., & McLaren, B.M. (2016). Learning with intelligent tutoring systems and worked examples: A study on ITS authoring tools. *User Modeling and User-Adapted Interaction: The Journal of Personalization Research*, 26(4), 339–384. <https://www.cs.cmu.edu/~bmclaren/pubs/NajarMitrovicMcLaren-LearningWithITSAndWorkedExamples-UMUAI2016.pdf>

⁹⁶ U.S. Department of Education, Office of Educational Technology. (2023, May). *Artificial intelligence and the future of teaching and learning: Insights and recommendations*. <https://www2.ed.gov/documents/ai-report/ai-report.pdf>

⁹⁷ Harrison, W., Brown, J., & Higgins, S. (2023, May). *A pilot impact study to evaluate the effectiveness of Eedi on raising attainment in mathematics at KS3 (11 years - 14 years old)*. What Worked Education. https://assets-global.website-files.com/5d11bcacde2cbd6f7d5d8955/649d83bd8b2660e3b5d0d575_Eedi%20Impact%20Pilot%20Evaluation%20May%202023.pdf

Kyron Learning is another AI-powered tutoring system whose designers chose not to incorporate generative AI to start. (Kyron Learning was one of the first 10 AI products reviewed by Common Sense Media,⁹⁸ and the head of Common Sense's social impact investment fund is one of the company's advisers.⁹⁹) Like Ben Caulfield, Kyron's CEO Rajen Sheth recognizes the power of great teachers, as well as the fact that not every student has access to a great teacher.

With Kyron Learning, teachers record video snippets, and the system uses natural language processing and speech recognition, which are both forms of AI, to understand and classify them. As a student moves through a lesson, the AI "listens" to them, determines whether they have misconceptions, and aims to serve the right video at the right moment.

Kyron first launched with fourth-grade math in a small cohort of schools for the 2023–24 academic year using conversational AI vs. generative AI—the former being a tighter, more guided system based on teachers' set of answers to specific use cases. It was also a way, says Sheth, to guard against errors, or so-called hallucinations, in generative AI. Student answers, stripped of personally identifiable information, are used only to train the model on classifying videos correctly—if a student answers X, the system should serve up video Y.

The company has since introduced generative AI into the product, but in a limited way. Teachers using the tool have the choice between the original approach, with a defined set of choices and responses, or having generative AI help guide the student to an answer. Sheth cautions that Kyron uses generative AI only to guide a student from a teacher-provided question to a teacher-provided answer. As a result, he says, it's safer than an open-ended chatbot, and teachers remain in control.

A fourth grade teacher who used Kyron Learning in her classroom in the 2023–24 school year says she appreciates the way the tool identifies and groups misconceptions. "That's something that I'll need to go back and reteach ... It's helping me make more clear, refined decisions." [See the Appendix for the [full Q&A](#).]

Khanmigo, from Khan Academy, is another product reviewed by Common Sense as part of its inaugural AI ratings system.¹⁰⁰ Of the three tools mentioned here, it's the only one that uses generative AI in a student interface, producing conversational step-by-step support and feedback in a chatbot format as students engage in independent learning and practice. Powered by OpenAI's GPT-4 and GPT 3.5, and further trained on Khan Academy's own learning content, it was one of the first generative AI learning tools to hit the market, in March 2023. The underlying models are what make the tool more expansive than our previous examples, and able to respond to many kinds of human input. That expansiveness, however, also means it's more prone to delivering incorrect information.¹⁰¹ Expansiveness costs money too—processors to run queries are expensive. Khanmigo cost \$4 per month, or \$44 per year, before the company entered into a partnership with Microsoft in May 2024. Microsoft Azure will donate access to its AI cloud infrastructure so that Khanmigo for Teachers can be offered free of charge to K–12 educators in the U.S., though parents and students will still have to pay a monthly charge.¹⁰²

⁹⁸ Common Sense Media. (n.d.). *AI initiative: Understanding the impact of AI on our kids*. <https://www.common Sense Media.org/ai>

⁹⁹ Kyron Learning. (n.d.). *Company page*. <https://www.kyronlearning.com/company>

¹⁰⁰ Common Sense Media. (n.d.). *AI initiative: Understanding the impact of AI on our kids*. <https://www.common Sense Media.org/ai>

¹⁰¹ *Ibid*.

¹⁰² Bishop, T. (2024, May 21). *Microsoft will make Khan Academy's Khanmigo AI tool free to U.S. teachers in cloud shift*. GeekWire. <https://www.geekwire.com/2024/microsoft-and-khan-academy-partner-on-khanmigo-ai-tool-for-teachers-in-cloud-shift/>

Creativity

Where adaptive learning tools often focus on closed tasks and individual progress, generative AI tools, with proper scaffolding and human oversight, have the potential to support students on open-ended creative tasks.¹⁰³ In New York City, Azadeh Jamalian runs what she calls a kids' co-design company. With a background in systems engineering and a PhD in cognitive science in education, Jamalian founded the Giant Room to create a learning organization that helps kids come up with ideas and engage in co-designing practices with designers, artists, and engineers. With generative AI, her approach has become easier to scale. AI has also made the Giant Room's programming more accessible. Jamalian has plans for a web platform that kids all over the country and eventually the world could use, at a much lower cost than her in-person workshops.

The way we're thinking about it is not necessarily replacing the instructors, not thinking of the AI as another person. How can we design this AI tool as another collaborator, as another tool in kids' creative process? — Azadeh Jamalian

"The way we're thinking about it is not necessarily replacing the instructors, not thinking of the AI as another person. How can we design this AI tool as another collaborator, as another tool in kids' creative process?" she says. Because generative AI uses natural language, it can be a way for kids to practice providing feedback and to build collaboration skills.

As an example, Jamalian offers a recent online product design bootcamp to help kids design their own valentine cards. On each day of the program, families received a creative prompt. On day one, kids were asked what "love monster" they wanted on the card, and they'd create one in whatever way they wanted, such as with crayons or with Legos, and send in a photo. The AI turned the photo into a digital avatar, and the kids would give feedback on it—how close was it to their vision, and how could it get even closer?

The next day was about how to craft the card's message, and the day after that an act of kindness. Finally, the kids chose their card's background color, and learned about the psychology of color as well as the cultural significance of different colors. Along the way, the kids talked with each other about their ideas, and, importantly, with adults, reinforcing the importance of having humans in the loop.

"They all went through the same design and thinking process, but every single output was really different. The AI is a tool that helps them learn," Jamalian says.

Jamalian has spent a lot of time thinking about how kids learn, and feels that too many conversations revolve around AI as something to learn *from*. That way of thinking about AI is dangerous, she says, because it can cause

¹⁰³ European Commission, Directorate-General for Education, Youth, Sport, and Culture. (2022). *Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators*. Publications Office of the European Union. <https://data.europa.eu/doi/10.2766/153756>; Mills, K., Ruiz, P., & Lee, K. (2024, February 21). *Revealing an AI literacy framework for learners and educators*. Digital Promise. <https://digitalpromise.org/2024/02/21/revealing-an-ai-literacy-framework-for-learners-and-educators/>

kids to lose confidence in their own intelligence. "All of a sudden there's this [AI] tutor that knows everything, and they can never be as smart as this. Whereas if you think of AI as a collaborator, you think about what are the things this tool can do, and what are you bringing in? And then you have the voice, the authority to give feedback." Kids learn to explain why they don't like an avatar created using generative AI, and to articulate their feedback in a way that the AI's next version is closer to what the child has in mind.

AI neither automatically enhances a child's creativity nor destroys it, says Jason Yip, PhD, associate professor at the Information School at the University of Washington, and an adjunct assistant professor in the university's Department of Human-Centered Design and Engineering. Yip has studied participatory design methods, and how technology can support parents, children, and neighborhoods learning together.

With generative AI, Yip says, children have to learn how to figure out the prompts and the tool's limitations. "They have to figure out what it can do, what it can't do. And they need adults at this point to guide them because it's not like magically the youth of America just know how to use generative AI."

Like Jamalian, Yip believes in pairing children and adults in co-design. One reason is that adults can usually draw or write faster—they're taking prompts from kids. Yip sees AI as similar to that role of reducing cognitive load. He mentions a story a child wanted to write about Marty McFly from *Back to the Future* but struggled drawing a DeLorean [sports car]. Instead of spending most of their time trying to draw a complex car, a young person can prompt the AI tool to draw it, and then concentrate on storytelling.

Yip's own research emphasizes that generative AI can't stand in for actual knowledge. Having an AI tool draw a DeLorean doesn't mean a child gains an understanding of color theory, perspective, or shading techniques. But if scaffolded correctly, generative AI can encourage a child to consider their larger creative intentions and experiences—it can help kids believe they're creative.¹⁰⁴

One other area in which generative AI is no substitute: actually assessing creativity.¹⁰⁵ Here again, human educators are needed.

The examples above all tie back to the model of "human in the loop" and bring this concept to life. Generative AI tools on their own may or may not help kids learn. Teacher oversight, in the form of thoughtful pedagogy and intentional learning design, is the essential piece.

Project-based learning and collaboration

Before Halloween last year, English language arts teacher Adam Aguilera, from Evergreen, Washington, came up with an idea for his eighth grade class: to write a story for the holiday in the style of a Choose Your Own Adventure book. (Note: Aguilera served as an adviser on this paper, and reviewed it in draft form.)

¹⁰⁴ Newman, M., Sun, K., Gasperina, I.D., Shin, G.Y. Pedraja, M., Kanchi, R., Song, M.B., Li, R., Lee, J.H., & Yip, J.C. (2024). "I want it to talk like Darth Vader": Helping children construct creative self-efficacy with generative AI. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. <https://dl.acm.org/doi/10.1145/3613904.3642492>

¹⁰⁵ European Commission, Directorate-General for Education, Youth, Sport, and Culture. (2022). *Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators*. Publications Office of the European Union. <https://data.europa.eu/doi/10.2766/153756>

The class used ChatGPT to help create characters and personality elements to start, and then to help build up the story. "AI would give the class a choice, like where the characters go, and they would debate this or that option. As we played through it, more students wanted to be part of the story as characters, so they showed up in the middle. There were some hilarious situations that the generator created, and it became a bonding experience for the class."

For Alex Kotran of aiEDU, project-based learning, whether or not it involves AI tools, is critical for developing skills that employers are looking for. One of aiEDU's exercises uses the "lawful good" meme¹⁰⁶ to structure a conversation on the uses and ethics of facial recognition. Students create a so-called meme board and identify gray areas, collaborating with and learning from each other as they also learn about AI's capabilities, limitations, and risks.

At the Giant Room, Jamalian partnered with a New York City elementary school on its annual international studies unit. As part of the program, students research one of that year's selected countries and make models that represent that country. Working with the Giant Room, kids used generative AI tools to design animal avatars native to the country and images to illustrate a travel guide. As they worked on the project, students also learned about AI biases and hallucinations. At the end, the teacher reported a 100% completion rate (usually about 80% of students complete this project) and higher overall effort and engagement.

Jamalian is also working with New York City elementary school teachers on a guide to using AI image generator tools. Still in development, the guide is being assembled based on survey results and co-design sessions. It will describe how AI image generation works, explain the risks involved with using generative AI tools, and offer teachers step-by-step support on piloting some of the Giant Room's storytelling and design bootcamps in their classrooms, with an emphasis on group ideation and sharing.

The above examples illustrate ways that educators came up with creative applications of generative AI tools and themes not only to teach about the usage and the implications of those tools, but to achieve specific learning goals around collaboration.

Teacher-supporting

One immediate, understandable reaction to generative AI's advances has been the fear that technology will replace human teachers. However, the opportunities we highlight here echo much of the academic literature and practice-oriented reports by positioning AI as augmenting and elevating human teaching¹⁰⁷ and enabling more informed decision-making by teachers.¹⁰⁸ They also all touch on the theme of capacity building for teachers, an important benefit as the educator workforce continues to suffer from burnout.¹⁰⁹

¹⁰⁶ *Alignment charts*. (Updated 2024, March). Know Your Meme. <https://knowyourmeme.com/memes/alignment-charts>; Tiffany, K. (2020, March 5). A chart to explain your entire worldview. *The Atlantic*.

<https://www.theatlantic.com/technology/archive/2020/03/alignment-chart-memes-moral-worldview-fantasy/607561/>

¹⁰⁷ Bundy, A. Preparing for the future of artificial intelligence. (2017). *AI & Society*, 32:285–7. As cited in Humble, N., Mozelius, P. (2022). The threat, hype, and promise of artificial intelligence in education. *Discover Artificial Intelligence* 2, 22.

<https://doi.org/10.1007/s44163-022-00039-z>

¹⁰⁸ Chaudhry, M.A., & Kazim, E. (2022). Artificial intelligence in education (AIEd): A high-level academic and industry note 2021. *AI and Ethics* 2:157–165 <https://doi.org/10.1007/s43681-021-00074-z>

¹⁰⁹ Walker, T. (2022, February 1). *Survey: Alarming number of educators may soon leave the profession*. *NEA Today*.

<https://www.nea.org/nea-today/all-news-articles/survey-alarming-number-educators-may-soon-leave-profession>

Experts we spoke with for this paper acknowledged that teachers, as well as districts and state-level systems, will have to adapt to AI. Educators can't necessarily plug a generative AI tool into a lesson plan they've been using for years and expect students to learn the same way. They need to ask themselves what they want students to learn, and then ask how they can get their students there, says Jason Yip of the University of Washington. In Yip's college-level informatics class, he used to have students take exams to memorize research facts and work out problems. But with generative AI's prevalence, rote memorization isn't as important. "Instead, I need them to understand that generative AI, while powerful in being able to create research protocols, is incredibly limited in understanding humans for research. So what I've done now is instead of exams, I give LLM outputs verbatim and have students write me arguments about the 'good' and 'bad' of these prompts," he explains.

Yip, a former high school teacher, acknowledges he has more freedom as a university professor than an educator who must follow district or state standards. He says all teachers, however, regardless of grade level, will need the autonomy to adapt curriculum to generative AI usage—otherwise "they are going to be at the mercy of those tools, rather than actually using those tools for positive purposes."

Yip offers a standard high school text—*Hamlet*—as an example. If the standard curriculum includes a traditional essay on Hamlet's relationship with Ophelia, a student can easily turn to ChatGPT to generate that essay. But if the teacher has the authority to make changes and assign an essay comparing Taylor Swift and Travis Kelce to Ophelia and Hamlet, generative AI "isn't a plagiarism machine anymore," Yip says.

The inverse is also true: AI can adapt to teachers. Zachary Pardos, PhD, associate professor of education at the Berkeley School of Education, says that next-generation technology should adapt to teachers' individual style and approach. "Instead of the teacher being the one to align to the curriculum of the edtech, the edtech can now align with the teacher," Pardos says. "They should actually feel more at ease with the right kind of technology."

Lesson planning and content differentiation

According to several recent surveys, lesson planning ranks among the highest AI use cases for teachers.¹¹⁰ A 2020 McKinsey report estimated that teachers spend 10.5 hours per week, out of a total of 50 hours, on preparation.¹¹¹ AI can help reduce that prep time, as well as offer teachers additional input into lesson planning.

In March 2024, Kyron Learning beta-launched its second product, Kyron Studio, a generative AI tool that enables educators to create their own lesson to supplement class instruction.¹¹² A teacher can input the title of the lesson, the specific audience they're targeting, and the learning objectives. Kyron Studio generates a default lesson that includes videos, which the teacher can then modify to fit their needs. Teachers can also upload their own documentation, such as PDFs and curriculum materials, to be incorporated into the AI-generated lesson, and record their own videos if they wish. (Anything uploaded remains private to the teacher, Kyron says.)

¹¹⁰ RAND American Educator Panels. (Administered 2023, fall). *American teacher panel, 2023 teacher omnibus: AI questions data file*. RAND Corporation. <https://bentobento.info/surveys/235>

Walton Family Foundation. (2023, March 1). *Teachers and students embrace ChatGPT for education*. <https://www.waltonfamilyfoundation.org/learning/teachers-and-students-embrace-chatgpt-for-education>; aiEDU. (N.d.) *Pulse survey report*. <https://www.aiedu.org/pulse>

¹¹¹ Bryant, J., Heitz, C., Sanghvi, S., & Wagle, D. (2020, January 14). *How artificial intelligence will impact K-12 teachers*. McKinsey & Company. <https://www.mckinsey.com/industries/education/our-insights/how-artificial-intelligence-will-impact-k-12-teachers>

¹¹² Sheth, R. (2024, March 21). *Announcing Kyron Studio*. Kyron Learning. <https://www.kyronlearning.com/post/announcing-kyron-studio>

As students move through the lesson, the AI responds to misconceptions and breaks down challenging questions. Afterward, teachers receive a report on their students' interactions to help them evaluate effectiveness and determine next steps.

Kyron's CEO says that the launch of Kyron Studio represented a careful decision for the company in terms of incorporating generative AI into the product. "I think, really, the reason for waiting was safety," says Rajen Sheth. "In my time at Google, I worked a lot on generative AI. And at the time, generative AI was not very good. This was two, three years ago. And I think it's still to the point where it's good at some things and not good at other things. The other thing I realized there was that the successful AI deployments that we had were ones where there was a human in the loop."

Sheth says that Kyron Studio can do in five minutes what would usually take four to five hours. "But the principles still apply, which is we want to have a teacher in the loop. We want the teacher not only to be able to correct the AI but put their personality into this. And we want to give the teacher the choice of how they're asking questions." The teacher has more than choice with the tool—they remain in control. No lesson is deployed to a student until a teacher has reviewed and published it.

Translation, language modification that adapts text to different reading levels, and IEP (individualized education program) components are other areas where AI could support teachers. The U.S. Department of Education report, for example, notes that AI might provide a voice assistant that allows a student with a visual difficulty to hear material and respond to it, or enable a group of students with hearing difficulties to present their project using American Sign Language (ASL) and use an ASL-to-spoken-English translation capability to voice it for other students.¹¹³ Amanda Bickerstaff of AI for Education says that special education teachers are the most overworked and AI represents an opportunity to lower their load. She cautions that while AI can help with planning and process, the technology should not be used on its own to determine an IEP goal.

Coaching and feedback

Similar to one-to-one tutoring, an instructional coach working with an individual teacher can be a transformative experience. It's also incredibly resource-intensive, and tends to be short-term. One potential area of opportunity: AI tools that process natural language and provide teachers with real-time, ongoing feedback on their instructional practices. This kind of feedback can highlight when a teacher inadvertently dominates discussion, for example, or whether certain students aren't speaking up. Previous studies have found that higher student participation and engagement improve learning outcomes.¹¹⁴

Dora Demszky, PhD, an assistant professor in education data science at the Graduate School of Education at Stanford University, has focused much of her research on using AI to provide teachers with automated feedback.

¹¹³ U.S. Department of Education, Office of Educational Technology. (2023, May). *Artificial intelligence and the future of teaching and learning: Insights and recommendations*. <https://www2.ed.gov/documents/ai-report/ai-report.pdf>

¹¹⁴ Sedova, K., Sedlacek, M., Svaricek, R., Majcik, M., Navratilova, J., Drexlerova, A., Kychler, J., & Salamounova, Z. (2019). Do those who talk more learn more? The relationship between student classroom talk and student achievement. *Learning and Instruction*, Volume 63. <https://doi.org/10.1016/j.learninstruc.2019.101217>; Reuell, P. (2019, September 4). *Study shows students in 'active learning' classrooms learn more than they think*. Harvard Gazette. <https://news.harvard.edu/gazette/story/2019/09/study-shows-that-students-learn-more-when-taking-part-in-classrooms-that-employ-active-learning-strategies/>; Center for Teaching and Learning. *Learning through discussion*. Columbia University. <https://ctl.columbia.edu/resources-and-technology/resources/learning-through-discussion/>

Together with colleagues, she began her work in 2018 and started using a pre-ChatGPT large language model in 2019.

Since then, they have built on top of language models like RoBERTa and GPT-4, tuning the models on expert-annotated data and prompts to improve their capabilities, and developing a tool called M-Powering Teachers (the *M* stands for machine). Only researchers covered within the Institutional Review Board protocol can access the data, and names and other identifying information are stripped out for analysis of classroom transcripts.

Demszky designed M-Powering Teachers primarily as a research tool, though it's open source and available to anyone in beta [a pre-release testing phase] at mpoweringteachers.stanford.edu. The tool measures the distribution of talk among the teacher and students, as well as the amount of time each person speaks, and highlights key talk moments.

Demszky sees the tool as a scalable way to provide feedback, in some cases working in conjunction with a coach, who can review and annotate transcripts and review data analysis along with the teacher. In tests in two online teaching settings, using randomized controlled trials, educators improved their "uptake" of student ideas, acknowledging, reiterating, and building on students' contributions.¹¹⁵ Further research in an in-person classroom setting is ongoing.

Demszky is also a data science and linguistics fellow with TeachFX, a consumer-facing AI feedback tool. A fourth grade math teacher using TeachFX in her Atlanta classroom says she used to get feedback on giving kids more talk time. "When I first started using TeachFX, it was showing me I was talking a lot. Since I started monitoring, I let the kids do the work more, let the kids come to the board more instead of doing the work for them," she says. She adds that a student once thanked her for making class fun that day, which TeachFX recorded, giving her the chance to get a little boost whenever she wants. "Now I can hear that forever," she says. [See the Appendix for the [full Q&A.](#)]

"I see the teacher as a very important role," says Demszky. Tools can improve instructional quality, she says, but not by replacing teachers. Tasks that can be automated can and should still be reviewed by a human, but that automation means that teachers can spend more quality time with students and build stronger relationships. Her idea of adding capacity means reducing burnout, not increasing workload. "I'm not in the efficiency boat, like 'now teachers should help 150 kids instead of 30 kids.'" She's aware of the risks, too. "There's a huge potential. But there's a real risk of using these models in the wrong ways and making inequities worse by creating unequal access to human teachers by replacing them with the [AI] model." Another risk that Demszky acknowledges is using the technology to surveil teachers and/or students in harmful ways. She says M-Powering Teachers generates non-evaluative, non-judgmental insights that require context to interpret, making them harder to use as a punitive measure. She's also put boundaries around who can access the tool's information to mitigate this risk further.

¹¹⁵ Spector, C. (2023, May 8). *Feedback from an AI-driven tool improves teaching, Stanford-led research finds*. Stanford Graduate School of Education. <https://ed.stanford.edu/news/feedback-ai-driven-tool-improves-teaching-stanford-led-research-finds>

Grading and assessments

As with lesson planning and instructional feedback, grading and assessments are another area of opportunity that still requires human oversight. People and computers don't always notice the same things in student writing.¹¹⁶ Generative AI tools can provide a first layer of feedback before teachers evaluate student assignments, for example, but can never be used as a replacement.

Adam Aguilera, the ELA teacher in Washington state, believes that by using AI grading tools, teachers can spend more time with students in conversation about their writing and books they are reading—important aspects of teaching that most teachers don't have enough time for now. Similarly, assessments can be scored for high-level feedback, which could free up time for teachers to provide more one-to-one support, such as conferring with students about projects.

Berkeley professor Zachary Pardos points to the potential for AI to generate "on the spot" assessments that target a specific area of uncertainty. Instead of a teacher choosing from preexisting options, they could create one tailored to that particular question.

Demszky also sees potential with assessments, and not just to save time for teachers. With multiple-choice assessments, such as those given at the end of an instructional unit or in accordance with state standards, teachers are looking at output rather than the process of learning. "While these assessments imply that the primary goal of learning is knowledge transfer, in our work we seek to go beyond that, helping teachers foster students' agency, independence, creativity, and curiosity," she says. If generative AI can analyze open-ended written answers holistically, the technology might allow for formative assessments, in which teachers can monitor student learning in order to provide ongoing feedback. These types of assessments provide an opportunity for educators to evaluate those softer skills and understand *how* their students are learning, not just *what* they're learning.¹¹⁷ One of Demszky's current projects focuses on helping teachers provide feedback on student writing that encourages students to revise their essay while also giving them agency and control over the process. Demszky notes that the teacher is always in the loop with this work in progress; generative AI never stands in for humans.

System-supporting

While much of the focus on AI's potential in education to date has revolved around direct support for teachers and students, schools, districts and state-level systems can also benefit.

Data interoperability

Data interoperability is the ability for different systems to share and translate data. The goal in the education arena: All the tools, systems, and platforms that a school or district uses should work together seamlessly. Erin Mote, who launched the data interoperability-focused Project Unicorn as part of InnovateEDU, offers an example from the travel industry: You check in for a flight with Delta, whether online or in person at the airport. When you get to security, the TSA official who's checking IDs and boarding passes already has the information from

¹¹⁶ U.S. Department of Education, Office of Educational Technology. (2023, May). *Artificial intelligence and the future of teaching and learning: Insights and recommendations*. <https://www2.ed.gov/documents/ai-report/ai-report.pdf>

¹¹⁷ Eberly Center, Carnegie Mellon University. (n.d.). *Formative vs. summative assessment*. <https://www.cmu.edu/teaching/assessment/basics/formative-summative.html>

Delta—and all the other airlines operating at that airport—allowing you and your fellow passengers to proceed toward your departure gates.

Generative AI isn't needed for systems to talk with each other; the technology to create data interoperability has existed for a while. But generative AI may bring other benefits, such as democratizing the capacity for data-driven decision-making, says Mote. She believes that foundational computer science skills are still useful, but that generative AI can speed up the process of dashboard creation and data visualization tremendously. That would allow more people to build expertise. At the moment, says Mote, it's hard to find people who have these sought-after skills and to pay them enough to stay in the education sector.

Jennifer Frentress, EdD, the superintendent of the San Carlos school district in California, agrees. Districts need coders, she says, and generative AI can level the playing field so districts that can't afford expert human coders can still translate and analyze data—and provide educators with easy-to-read, actionable data. There's a risk of errors with AI, but Frentress doesn't see it as any more of a risk than with humans doing all the work. "It's crazy we are making high-stakes decisions with people who don't understand delimited structure," she says, meaning many districts often make decisions based on data without the expertise necessary to structure and interpret that data. (That's not to say there aren't critical questions around data and AI—we address those below in the section on AI's risks and challenges.)

Examples of how data pulled and analyzed from different tools and platforms could support system-wide planning and insights include resource allocation, scheduling, and diagnosing learning difficulties.¹¹⁸

Dan Carroll also sees areas of potential opportunity. He co-founded the company Clever precisely to aid schools with data interoperability. The platform provides a unified way to access multiple tools. It also ensures that any one tool can both bring in relevant data from other tools, and push its own data out across the system. Clever doesn't currently use generative AI, but Carroll sees possibilities for the technology in three areas: building first iterations of data integrations across different systems, creating early versions of data analysis (especially when the data is messy), and cross-state standards alignment.

Administration and logistics

Another area of potential opportunity is around administration and logistics. Jean-Claude Brizard of Digital Promise wants to see more of a spotlight on these areas. "What could AI mean for the back-office operations, transportation, food services, the stuff we know tend to be endpoints?" he asks.

Parent engagement

Parents often get left out of the conversation around opportunities for AI in education. But they represent a critical constituency. Better data interoperability and data analysis could mean that parents get a clearer picture of their children's progress. Schools could also use generative AI to communicate more frequently and more effectively with parents. Laurence Holt, senior adviser at XQ Institute, thinks an AI-powered chatbot is low-hanging fruit for districts. He sees potential for better explanations for families who have a hard time making sense of the data that the school shares with them. Parents could also more easily learn about academic and

¹¹⁸ European Commission, Directorate-General for Education, Youth, Sport, and Culture. (2022). *Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators*. Publications Office of the European Union. <https://data.europa.eu/doi/10.2766/153756>

extracurricular programs their kids may be eligible for and interested in, he says. As with any tool, though, districts need to approach with caution. The Los Angeles superintendent gave a splashy presentation in April 2024 to announce a chatbot for parents and students with a \$6 million price tag. Just two months later, the company behind the chatbot collapsed.¹¹⁹

Beyond providing parents with easily accessible information about their children's school experience, there's also an opportunity for schools to offer broader trainings for parents and caregivers on how to use generative AI themselves and with their families. These kinds of literacy offerings serve the dual purpose of increased support for students' education at home, and direct access to AI literacy programs for adults, especially important for lower-wealth families, immigrant communities, and families of color.

¹¹⁹ Goldstein, D. (2024, July 1). AI 'friend' for public school students falls flat. *New York Times*. <https://www.nytimes.com/2024/07/01/us/ai-chatbot-los-angeles-schools.html>

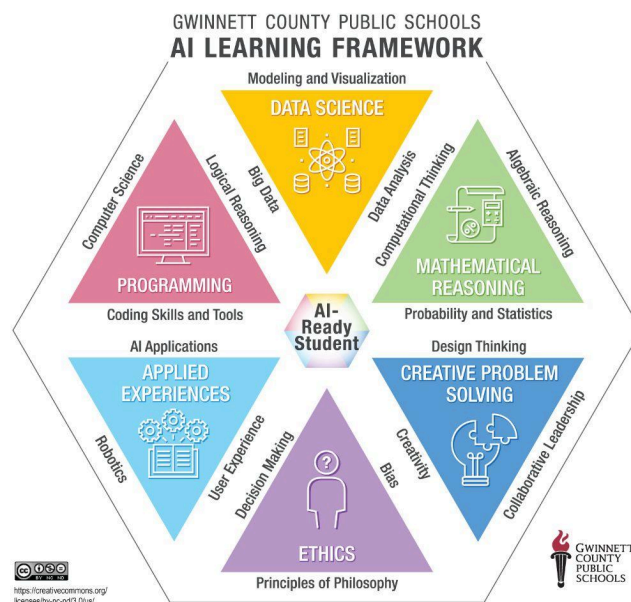
Seckinger High: All In on AI

Here's how one high school in Georgia has embraced AI as technology that all students need to learn about, and as technology to help students learn.

Right before ChatGPT launched in the fall of 2022, the first AI-themed high school in the country opened its doors. Now in its third year of operation, with about 2,300 students, Seckinger High School represents the anchor of Gwinnett County Public Schools' efforts to embed an AI curriculum across grades throughout a cluster of K-12 schools in the district.

Early grades start by building an ethical understanding of AI and learning how to empathize with others through design thinking. In middle school, students go deeper into algorithms, the foundations of coding, and data sets, along with conversations on ethical issues like data privacy and algorithmic privacy. By the time students hit high school and start actually using more AI-powered tools in earnest, they have "a critical eye," says Babak Mostaghimi, former assistant superintendent. "They can more broadly think about the purposes or use cases to solve problems with AI."

Gwinnett County uses a six-theme "learning framework" that includes ethics, applied experiences, creative problem-solving, mathematical reasoning, data science, and programming. Each theme is broken down into pre-K, K-2, 3-5, 6-8, and 9-12 grade levels in terms of expectations and requirements. The goal for teachers, according to Seckinger principal Jimmy Fisher, is to work at least one of these themes into each lesson. (Fisher acknowledges that not every single class in every single period will explicitly touch on a theme.)



The district also likes to use an aquatic analogy: swimming, snorkeling, and scuba diving. "Everyone who graduates from our schools needs to be able to swim at a minimum," Mostaghimi says. He likens AI "swimming" to basic technological capabilities, such as using a computer, joining a Zoom meeting, and, now, using some generative AI tools. The goal is to gain a foundational understanding of how AI works and learn how to avoid getting fooled by scams and misinformation.

"Snorkelers," he says, understand the tools more deeply. "How do I manipulate the tools to get what I need? How do I understand that what's coming out may be biased or influenced in different ways?"

Scuba divers are the equivalent of a student on a computer science track. The district wanted to invite more people to consider a path into the tech industry, especially girls and students of color, who have historically been underrepresented in the field. (About two-thirds of the 180,000 students in the district are Black or Hispanic/Latino.¹²⁰) Gwinnett County built a three-course career and technical pathway for the scuba divers. "But even that is focused on the technical tools and the creative problem-solving pieces and the ethical pieces, so we can have ethical creators," Mostaghimi says.

Seckinger and the district extend their approach beyond the classroom through local internships. Community partners on both a school advisory board and one at the district level share what they're looking for in employees, which has helped Seckinger determine how best to prepare their career and technical pathway students. Fisher says Seckinger is committed to identifying opportunities for "authentic" internship placements; more than 100 students participate in work-based learning.

As you might guess at an AI-focused high school, Seckinger never considered banning ChatGPT. Mostaghimi says that the school considered it everyone's responsibility to determine how to use the tool "carefully and ethically." As such, the school began discussions about how generative AI should and shouldn't be used. "You can't prevent a kid from using generative AI, so let's change the assignments we're giving. If you can just use ChatGPT to get an 'A' on an assignment, it means that wasn't a good assignment," Mostaghimi says about an early thread of discussion.

Principal Fisher says ChatGPT can help students in English language arts classes who might be struggling with essay writing. A student who normally can only eke out four sentences and instead hands in five paragraphs generated by ChatGPT will get a failing grade. "But if they can use that as a starting point, it just raises the bar for where they can go. [We're] trying to get our teachers to wrap their heads around that and use this as a stepping stool, as a resource, as a leap pad. And we have to understand that there is ethics involved in that. When we do any sort of research, we have to cite it the proper way."

Fisher says that the school has allowed teachers to experiment with different tools and platforms, while also providing them with learning opportunities led by a core group of teachers. The most popular tools are those built on large language models: ChatGPT, Gemini, and Claude. (Not all of Seckinger's teachers are AI masters just yet; Fisher says one upcoming learning opportunity will focus on the basics for those

¹²⁰ Gwinnett County Public Schools. *GCPS by the numbers*. <https://schools.gcpsk12.org/Accolades>

who "haven't jumped on the boat yet.") Each week, a newsletter goes out with suggestions of new tools and tips. When Fisher spoke with Common Sense in February, more formal acceptable-use policies and age-appropriate guardrails were in the works. The district announced in April that it was partnering with EdSAFE AI Alliance to create policy recommendations and educational materials for teachers, students, and parents.¹²¹

During Seckinger's planning stages, district leadership sought parent input. One of the biggest takeaways, says Mostaghimi, was that parents wanted their children to have a "normal high school experience" with opportunities to participate in music, arts, and sports, all of which Seckinger does provide. He says one change based on parent feedback was how they talked about the school. "Instead of talking about the technology, we started talking about preparing kids for the future."

Fisher sees parent education as an opportunity. He wants to offer more chances for families to come to the school for showcases of student work and to learn more about the technologies behind them. For Fisher, AI literacy is essential for parents, so they can help prepare their kids for the future as well.

Risks, harms, and challenges

Although generative AI presents many opportunities in a school setting, there are also substantive challenges. The risks detailed below are not exclusive to education, though we provide education-related examples. These risks are also applicable to all sectors and age groups. (Adults should know how to think critically and appropriately cite sources of information, for example.) What is specific to education is the fact that young people are at the heart of it. And as we noted in the introduction to this paper, young people demand special care and consideration from us as a society.

Our selection of risks, harms, and challenges, like the opportunities we highlight above, is not intended to be comprehensive or exhaustive. We left off the environmental and financial costs of building and running these massive models, as well as the implications regarding the further concentration of wealth and power in the hands of a few large tech companies, for example. These issues will no doubt translate into future costs for administrators, and in the case of environmental impact, moral objections (or at the very least worthy debate) from students and educators over generative AI's toll.

We also do not address nudification, AI-generated pornography, bullying, mental health, and other risks, though these issues will also bring their share of trouble to school leaders' doorsteps and affect young people in negative ways. Instead, we try to focus on the risks most directly applicable to education and learning environments, and the ones that those working in the space, be they educators, students, researchers, advocates, or policymakers, may have the best shot at influencing.

¹²¹ Gwinnett County Public Schools. (2024, March 29). *Gwinnett county public schools launches AI policy lab*. <https://www.gcpsk12.org/news-details/~board/news-releases/post/gwinnett-county-public-schools-launches-ai-policy-lab>

We also tried to reflect what we heard in our interviews for this paper, and what we saw from surveys in terms of teacher concerns. Top concerns that emerged from aiEDU and Department of Education surveys and listening sessions include:

- Lack of human interaction
- Increased cases of cheating
- Privacy issues
- Misinformation
- Dependence on technology
- Equity in access to technology¹²²
- Algorithmic transparency and user control of data
- Bias and lack of diversity¹²³

Inaccuracies and hallucinations

AI carries the risk of creating false information (also known as hallucinations) and repeating inaccurate information found within the sources used to train the system.¹²⁴ AI literacy can mitigate this risk by making people aware of this tendency, and by teaching them to fact-check. Certain products have restricted their training data to only vetted sources to reduce this risk further. But this is still a huge problem in education, where students in particular may trust multipurpose generative AI tools like ChatGPT to provide trusted, factual information. Professor Pardos at Berkeley wonders if and when large language models will hit a performance reliability level that other kinds of technologies have, such as machine learning. "Where is the plateau? Is it close to the performance we are at today? Or is it going to double?" he asks. "That's an open question: Are we actually going to get to an error rate that is usable in all domains at the K-12 level? That's not clear yet."

Bias and inequity

Just as generative AI can repeat misinformation as fact, it can churn out bias as fact. Common Sense's own resource on how to handle AI in schools explains it well: "AI can only learn from its source(s), so it takes on the biases, misinformation, and problematic content of the original material. And if the team of developers isn't representative, it's almost guaranteed that implicit bias will be woven into the framework of the tool."¹²⁵ When ChatGPT was asked to speak in the cadence of the author of a book featuring a Black protagonist, for example, the tool added the word "yo" in front of random sentences.¹²⁶ As more people use these tools, more people are exposed to and can spread problematic content, leading to the risk of amplification of preexisting biases.¹²⁷

It's not enough to try to catch bias after it's out there. A recent episode of the *Sunday Show* podcast from Tech Policy Press featured University of Washington professor Batya Friedman, PhD, who pioneered value sensitive design, an approach to account for human values in the design of technical systems.¹²⁸ On the podcast, she called

¹²² aiEDU. (N.d.) *Pulse survey report*. <https://www.aiedu.org/pulse>

¹²³ U.S. Department of Education, Office of Educational Technology. (2023, May). *Artificial intelligence and the future of teaching and learning: Insights and recommendations*. <https://www2.ed.gov/documents/ai-report/ai-report.pdf>

¹²⁴ Ibid.

¹²⁵ Elgersma, C. (2024, March 6). *ChatGPT and beyond: How to handle AI in schools*. Common Sense Media. <https://www.commonsense.org/education/articles/chatgpt-and-beyond-how-to-handle-ai-in-schools>

¹²⁶ Chen, C. (2023, March 9). *AI will transform teaching and learning. Let's get it right*. Stanford University Human-Centered Artificial Intelligence. <https://hai.stanford.edu/news/ai-will-transform-teaching-and-learning-lets-get-it-right>

¹²⁷ MIT Sloan Teaching & Learning Technologies. (n.d.). *When AI gets it wrong: Addressing AI hallucinations and bias*. <https://mitsloanedtech.mit.edu/ai/basics/addressing-ai-hallucinations-and-bias/>

¹²⁸ Information School, University of Washington. [n.d.] *Batya Friedman, faculty profile page*. <https://ischool.uw.edu/people/faculty/profile/batya>

for investing resources in developing better methods for identifying or anticipating biases early on in generative AI systems, similar to the way engineers have developed standards for buildings to withstand earthquakes, and improve them over time.

Friedman also called for new techniques of reporting bias after the fact. "Those techniques and processes need to take into account not just those people who are technically savvy and who have access to technology, but to recognize that people who may never put their hands on the technology might be significantly affected by biases in the system," Friedman said on the podcast. "And that they need ways, that are perhaps not technical ways, of communicating harms they are experiencing and then having those being taken up seriously and addressed. I would say that we're at the very beginning of learning how to do that."¹²⁹

Lack of representation

The massive data sets that generative AI tools are built on also bear another form of bias: a lack of inclusion and representation. As the song from *Hamilton* goes, "Who tells your story?" If you're relying on generative AI, the answer may sometimes be "no one."

In some cases, it's a problem of language itself. The pivotal Stochastic Parrots paper, citing work by other researchers, noted that "over 90% of the world's languages used by more than a billion people currently have little to no support in terms of language technology."¹³⁰

Another paper synthesizing research on the risks of AI foundation models (general purpose systems such as those that power ChatGPT and similar chatbots¹³¹) in education said that these models "may require speakers of minoritized varieties to accommodate to dominant varieties in educational contexts, incurring higher costs for these speakers and denying them the legitimacy and use of their varieties."¹³² The paper pointed to AI-automated feedback and grading as an area of particular concern. If the models are trained on biased data to begin with, those biases can seep into the feedback, which in turn "may induce students to adopt writing styles that mirror dominant cultures."¹³³ The end result of this feedback loop can be another form of bias amplification, where foundational models may reproduce harmful assumptions about how students should be taught and the legitimacy of certain languages and language varieties.¹³⁴

Information censorship

A recent article from *The Markup* detailed how some school districts block certain content via web-filtering software. Students were prevented from accessing the Trevor Project, a suicide prevention and crisis intervention nonprofit for LGBTQ+ young people; reproductive health information; a video of the artist Picasso in the process

¹²⁹ Hendrix, J. (2023, December 17). *What are we building, and why?* Tech Policy Press.

<https://www.techpolicy.press/what-are-we-building-and-why/>

¹³⁰ The "**Stochastic Parrots**" paper is making reference to Joshi, P., Santy, S., Budhiraja, A., Bali, K., & Choudhury, M. 2020. The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 6282–6293. <https://doi.org/10.18653/v1/2020.acl-main.560>

¹³¹ Jones, E. (2023, July 17). *Explainer: What is a foundation model?* Ada Lovelace Institute.

<https://www.adalovelaceinstitute.org/resource/foundation-models-explainer/>

¹³² Blodgett, S.L., & Madaio, M. (2021, October 19). *Risks of AI foundation models in education*. <https://arxiv.org/abs/2110.10024>, referring to Baker-Bell, A. (2020). *Linguistic justice: Black language, literacy, identity, and pedagogy*. Routledge.

¹³³ Ibid, referring to Elijah Mayfield. Individual fairness in automated essay scoring. In *Proceedings of the Workshop on Contestability in Algorithmic Systems (CSCW)*, 2019.

¹³⁴ Ibid.

of painting; and analyses of *The Odyssey*, to name just a few examples.¹³⁵ It's easy to imagine ways that generative AI models could also be tuned to prevent access to certain information, especially since education has long been a space where decisions follow political alignments. Kids who don't have access to non-school devices, and/or who don't have reliable internet access outside of school, would be impacted the most—another potential instance of inequity.

Cheating and plagiarism

When ChatGPT was released in late 2022, many schools and districts immediately banned it outright, due to concerns about academic integrity and inaccurate information. Some reversed course soon after,¹³⁶ though restrictions often remain in places where the tool is allowed. Further, as we discussed earlier in this paper, most school systems still lack specific guidelines on generative AI usage, and there's no accepted, detailed definition of generative AI and cheating.

Our conversations with teens and parents revealed a patchwork system that varies by teacher and by school, with some encouraging the use of generative AI tools in specific instances, and others continuing to enforce a ban and threaten students with AI detection tools. No matter where a school falls along the spectrum, though, no one seems to have clearly answered the question: Where's the line between acceptable use of generative AI and cheating or plagiarism?

While it may take time to formulate consensus, the use of AI detection tools represents an even more urgent problem. A 2023 Center for Democracy and Technology survey of educators reported an increase from the 2022–23 school year in both the use of AI content detection tools and student disciplinary action¹³⁷—cause for concern, given that research shows these tools do not always accurately differentiate between AI-generated and human-written text.¹³⁸ The same survey also found that the majority of teachers have not received guidance on responding to a suspicion of generative AI usage against school rules.¹³⁹

As the survey's writeup points out, disciplinary consequences also present greater risks to certain groups of students. Nearly half of teachers agree that students who use school-provided devices are more likely to get in trouble or face negative consequences for using generative AI. And previous Center for Democracy and Technology research has shown that Black, Hispanic, rural, and lower-income students rely more heavily on school-issued devices.¹⁴⁰

¹³⁵ Mathewson, T.G. (2024, April 13). *Schools were just supposed to block porn. Instead they sabotaged homework and censored suicide prevention sites.* The Markup.

<https://themarkup.org/digital-book-banning/2024/04/13/schools-were-just-supposed-to-block-porn-instead-they-sabotaged-homework-and-censored-suicide-prevention-sites>

¹³⁶ Banks, D.C. (2023, May 18). *ChatGPT caught NYC schools off guard. Now, we're determined to embrace its potential.* Chalkbeat New York.

<https://www.chalkbeat.org/newyork/2023/5/18/23727942/chatgpt-nyc-schools-david-banks/>

¹³⁷ Dwyer, M., & Laird, E. (2024, March). *Up in the air: Educators juggling the potential of generative AI with detection, discipline, and distrust.*

Center for Democracy & Technology.

<https://cdt.org/wp-content/uploads/2024/03/2024-03-21-CDT-Civic-Tech-Generative-AI-Survey-Research-final.pdf>

¹³⁸ [Ibid.](#), citing Elkhatat, A.M., Elsaid, K., & Almeer, S. (2023, September 1). Evaluating the efficacy of AI content detection tools in differentiating between human and AI-generated text. *International Journal for Educational Integrity*. perma.cc/HS8H-7RL6.

¹³⁹ [Ibid.](#)

¹⁴⁰ [Ibid.](#), citing Laird, E., Grant-Chapman, H., Venzke, C., & Quay-de la Vallee, H. (2022, August 3). *Hidden harms: The misleading promise of monitoring students online.* Center for Democracy & Technology. perma.cc/E3E2-CE2Q.

<https://cdt.org/wp-content/uploads/2024/03/2024-03-21-CDT-Civic-Tech-Generative-AI-Survey-Research-final.pdf>

The survey also found that disciplinary actions for generative AI use increase at schools that ban the technology in classrooms.¹⁴¹ And that could present yet another problem of equity. "AI is going to show up in every edtech product if it isn't already there. You can't ban it unless you ban all tech in your classrooms," says Digital Promise's Jean-Claude Brizard. "And I worry because [bans] tend to be in poorer districts where kids are not going to get access to this technology."

Overreliance on technology and loss of critical thinking

In an article in *The New Yorker* magazine, the science fiction writer Ted Chiang wrote, "If there is any lesson that we should take from stories about genies granting wishes, it's that the desire to get something without effort is the real problem."¹⁴² Likewise, the fourth grade teacher in Los Angeles County we spoke with shared her fear that AI is erasing "productive struggle." Amanda Bickerstaff of AI for Education says she hears this concern a lot. One of the biggest questions educators are grappling with, she says, is, "Are kids going to cognitively offload so much that they stop learning?" I think this is where we want to be very cautious about the ways in which we approach this, because that [concern] is real."

The calculator example came up again and again in our interviews and research for this paper. Jean-Claude Brizard of Digital Promise, who taught physics in New York City earlier in his career, pointed to New York state's allowance of graphing calculators on state regents physics exams. It became an accepted tool, he says, when they included the directive to students to show their work. Similarly, Azadeh Jamalian of the Giant Room sees generative AI as a tool, like a calculator, that engages with you to find the answer. And an eighth grader in northern California offers a similar take: "Just because we have calculators, you still need to learn long division, because there's a whole thought process to it. So it's really about, when's the right time to use [AI tools]. It's about minimizing the workload, but then still maximizing the education that you're getting."

Just because we have calculators, you still need to learn long division, because there's a whole thought process to it. So it's really about, when's the right time to use [AI tools]. It's about minimizing the workload, but then still maximizing the education that you're getting. – eighth grade student

Generative AI is much more powerful than a calculator, though—it can touch and impact every subject taught in school. At a 2023 Stanford summit on AI and education, one political science professor noted that writing is a way of learning how to think, and asked whether outsourcing much of that work to AI would harm students' critical thinking development.¹⁴³ Beyond its impact on critical thinking, Eisha Buch, head of teaching and learning at

¹⁴¹ Dwyer, M., & Laird, E. (2024, March). *Up in the air: Educators juggling the potential of generative AI with detection, discipline, and distrust*. Center for Democracy & Technology.

<https://cdt.org/wp-content/uploads/2024/03/2024-03-21-CDT-Civic-Tech-Generative-AI-Survey-Research-final.pdf>

¹⁴² Chiang, T. (2023, May 4). Will A.I. become the new McKinsey? *New Yorker*.

<https://www.newyorker.com/science/annals-of-artificial-intelligence/will-ai-become-the-new-mckinsey>

¹⁴³ Chen, C. (2023, March 9). *AI will transform teaching and learning. Let's get it right*. Stanford University Human-Centered Artificial Intelligence. <https://hai.stanford.edu/news/ai-will-transform-teaching-and-learning-lets-get-it-right>

Common Sense, suggests that generative AI is transformative enough that we'll need to, as a society, communicate "why this is important. Why is that thought process actually important now? Or are there different thought processes we need to help cultivate?"

Data security

Consider the vast amount of data in education already: personally identifiable information such as name, age, and home address; class attendance; grades; and test scores, to name the most obvious. Then, layer in additional data from AI-powered platforms and tools: diagnostic assessments, misconceptions, evaluations of competencies, and inferred emotional states. Who owns the data, who is able to access it, and who is held responsible if something goes wrong?¹⁴⁴ Since 2016, the K12 Security Information Exchange has cataloged a total of 1,331 publicly disclosed school cyber incidents affecting U.S. school districts and other public educational organizations. The nonprofit group estimates that, because of weak disclosure requirements, the true number may be up 10 to 20 times higher.¹⁴⁵ Their 2022 report warned that "the absence of meaningful cybersecurity risk management standards for schools at either the state or federal levels—coupled with a lack of resources dedicated to meeting any such standards—all but guarantees that many districts will continue to place the safety and security of students, teachers, and community members at avoidable risk."¹⁴⁶

The Institute for Ethical AI in Education identifies privacy and data governance—ensuring that everyone has the right to control their data—as a requirement to mitigate risks for learners.¹⁴⁷ Roxana Marachi, PhD, professor of education at San Jose State University, has questioned the amount of control that students and parents actually have over data. In an interview with the news outlet *The Markup*, she highlighted the "right to be forgotten" as a key principle. "We need to consider what consequences will result when digital labels follow students throughout their educational paths," she said.¹⁴⁸ Likewise, a 2024 report from the National Education Policy Center questions whether students and parents will have ways to challenge judgments and decisions by AI-based applications, such as particular academic pathways or advancements.¹⁴⁹

David Ball, who leads the Headstream initiative on adolescent well-being, says he's been hearing more from young people lately about data concerns. In Headstream's most recent selection process for their accelerator program, which centers young people in building solutions for mental health, applicants raised questions about data and trust. "Why should I trust a particular product, approach, or methodology? If I open up to a mental health platform, and provide information, what's their incentive? What are they doing with my information?" A March 2024 Edelman analysis found similar trust issues. Their Trust Barometer found that globally, trust in AI companies

¹⁴⁴ Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Buckingham Shum, S., Santos, O.C., Rodrigo, M.T., Cukurova, M., Bittencourt, I.I., & Koedinger, K.R. (2022). Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, 32:504–526. <https://doi.org/10.1007/s40593-021-00239-1>

¹⁴⁵ K12 Security Information Exchange. (2022). *The state of K-12 cybersecurity: Year in review*. <https://static1.squarespace.com/static/5e441b46adfb340b05008fe7/t/6228bfe3f412c818293e16e1/1646837732368/StateofK12Cybersecurity2022.pdf>

¹⁴⁶ Ibid.

¹⁴⁷ Chaudhry, M.A., & Kazim, E. (2022). Artificial intelligence in education (AIED): A high-level academic and industry note 2021. *AI and Ethics* 2:157–165 <https://doi.org/10.1007/s43681-021-00074-z>

¹⁴⁸ Angwin, J. (2022, January 15). *The big business of tracking and profiling students*. Newsletter: Hello World. The Markup. <https://themarkup.org/newsletter/hello-world/the-big-business-of-tracking-and-profiling-students>

¹⁴⁹ Williamson, B., Molnar, A., & Boninger, F. (2024, March). *Time for a pause: Without effective public oversight, AI in schools will do more harm than good*. National Education Policy Center, School of Education, University of Colorado Boulder. https://nepc.colorado.edu/sites/default/files/publications/PB%20Williamson_0.pdf

has declined over the past five years from 61% to 53%. In the U.S. over the same period, trust is even lower, dropping from 50% to 35%.

Consent and privacy

Consent and privacy are two other issues of concern regarding data. Many schools send a consent form to parents at the start of the school year regarding the use of technology, but they're generally not very specific. Seckinger High School in Georgia, which explicitly weaves AI throughout all subjects, tells parents that generative AI is "a tool that teachers may use to help support classroom instruction that is relevant to the future of students" and that these AI tools "fall under the same guidelines of appropriate use of technology as other technologies" used in the district. The consent form states that "generative AI tools have features that may require us to collect, use, and/or disclose information that your child produces" but provides no example other than a prompt that a student might input into a chatbot.¹⁵⁰

What gets typed into a chatbot is critical, because this data is often used to retrain generative AI models—and that carries privacy implications. Unless a user opts out of this baked-in feature (it's usually difficult to figure out if this option exists, and if so, how to activate it), the information a user puts in is now part of the system. That includes any personally identifiable or sensitive information—a huge risk for children, who may not know what they should and shouldn't disclose. In addition to retraining the models, sensitive information can be used to build a unique profile of users over time, says Common Sense's head of privacy Girard Kelly. "Companies can use that information for their own manipulative or persuasive purposes, such as tailoring AI product answers with blurred advertising to shape opinions and preferences for brands over time. In the future, you may get customized responses to your queries that are designed for the benefit of the AI product or advertisers, not the user," he says.¹⁵¹

Indeed, many people may give their consent for data collection without knowing the extent of the information they are sharing, or how it will be used.¹⁵² A 2016 research briefing on student privacy and edtech (well before generative AI entered the educational scene) noted that teachers may adopt certain tools at the classroom level, which may result in circumventing consent requirements, even if unwittingly.¹⁵³ Further, companies may integrate AI features into products already on the market without any additional notice or consent protocols for the new purposes of data collection and usage, says Kelly. "Parents can't provide informed consent if they don't have a meaningful choice."

Further complicating matters, different tools have different requirements regarding parental consent. ChatGPT, for example, restricts usage for children age 12 and under, and requires parental consent for kids age 13 to 17. But the former is easy to circumvent (because ChatGPT no longer requires users to log in¹⁵⁴), and the latter isn't checked or enforced.

¹⁵⁰ Seckinger High School AI consent form, as provided to Common Sense Media via email.

¹⁵¹ Federal Trade Commission. (2022). Protecting kids from stealth advertising in digital media, testimony of Girard Kelly. https://www.ftc.gov/system/files/ftc_gov/pdf/FTC-Protecting-Kids-from-Stealth-Advertising-in-Digital-Media%E2%80%93October-19-2022.pdf

¹⁵² Akgun, S., & Greenhow, C. (2022). Artificial intelligence in education: Addressing ethical challenges in K–12 settings. *AI and Ethics*. 2:431–440. <https://doi.org/10.1007/s43681-021-00096-7>

¹⁵³ Plunkett, L., & Gasser, U. (2016, September 28). *Student privacy and ed tech (K–12) research briefing*. Networked Policy Series, Berkman Center Research Publication No. 2016-15. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2842800

¹⁵⁴ David, E. (2024, April 1). *Now you can use ChatGPT without an account*. The Verge.

<https://www.theverge.com/2024/4/1/24118007/openai-chatgpt-free-no-account-login>

Researchers have questioned whether consenting to use AI-powered tools in the classroom is actually a choice at all if these systems are required by public schools.¹⁵⁵ Roderic Crooks, PhD, assistant professor in the department of informatics at the University of California at Irvine, calls it "the illusion of an opt-out." When you introduce a platform into a public service, he says, you're essentially requiring people to use it. Crooks sees danger in the amount of data that tech companies have been ingesting for years, even before what he terms the "current mania for AI." The advent of generative AI, he believes, represents an intensification of attempts to "collect as much data as possible, to convert that data into an asset, and then to profit from that asset."

In his research, Crooks has found that students and classroom teachers don't know much about what kinds of data are being collected and where it's going. In his experience, he says, principals and administrators might develop some expertise in data for limited, directed purposes. Then, up a rung, a district leader might be responsible for negotiating contracts with platforms. "Essentially, the higher up you go, the more people will know about data and where it goes," Crooks says. He sees parallels to tech companies benefiting financially. "The benefits of data, reputational and financial, accumulate at the top. The benefits can only go where people are positioned by wealth and authority to reap those benefits," he says.

Human replacement

Even if we agree with the critiques mentioned earlier of "longtermism" (the focus on protecting the future of humankind while downplaying immediate, potentially solvable problems with AI), there are valid concerns today about what we should outsource to technology and what should remain part of the human domain. As an example, Jason Yip of the University of Washington brings up advances in tools for teen mental health and asks whether we're trying to scale where we shouldn't be. He has published research on AI's inability to "code switch," meaning that the system doesn't know whether it's interacting with a child or an adult, a native English speaker or a recent immigrant. He sees this as a problem with mental health use cases for AI. "I get that therapists are hard to find, but that ability to code switch is so important. If a teen is experiencing suicide ideation, I want them on a phone with a hotline as fast as possible. I don't want them with a generative AI tool. We have to focus on what the limits are," he says.

David Ball from Headstream says that recent accelerator applicants consistently said they want to talk to a real person, whether in person or by video or phone. But he also explains, "I've seen evidence and data that conflicts with that," offering a hypothesis that if a young person is dealing with stigma around mental health issues, or they're exploring their own feelings early on in the process, a chatbot may be an appropriate solution. AI might also be used in intakes and assessments. But he says that he and his team take the desire for human personalization seriously. "What is the efficacy of using AI when it's actually replacing the human element?" he wonders.

Amara, the AI ethics researcher, recalls a workshop with pastors and divinity school students. She says they noticed the algorithm they were exploring in the session interpreted scripture in a particular way, or used a version of the Bible they don't usually rely on. Amara says one of her takeaways was that you can't turn to these tools for spiritual concerns, because it's so personal. "What does it mean to go to a chat box for something that a human is supposed to give?" she asks.

¹⁵⁵ Ibid.

What students think

Much of this landscape review has focused on data, ideas, and policy from adults—but young people are a critical group of users of AI tools in an educational setting, and it is critical that we ask them about their use of these products. As we noted earlier, a November 2023 Common Sense survey found that half of young people age 14 to 22 have used generative AI at some point in their lives, and according to Impact Research, about half of adolescents age 12 to 18 use ChatGPT at least once a week, both in and out of school. While survey results are informative, we also wanted to talk directly with young people to get a qualitative sense of their feelings and beliefs.

We spoke with several students in different parts of the country, including four from the Common Sense Youth Research Council, which provides input into research questions and data analysis around children and technology. Several themes emerged from our conversations:

For most kids, ChatGPT equals AI

When we asked students what comes to mind when they hear the term "artificial intelligence," ChatGPT topped the list. The OpenAI chatbot has become nearly synonymous with AI. Other AI-powered or AI-enhanced tools that students mentioned included image generators, design programs, the learning tool Quizlet, the writing assistance app Grammarly, and the language learning app Duolingo.

AI policies vary among districts, schools, and teachers

Whether and how kids can use AI for schoolwork varies widely, and it often depends on how familiar an individual teacher is with the technology. The lack of consistency is confusing, and kids agreed that educators need more training in order to develop clearer guidelines.

Students don't want bans, but they do want clear guidelines

Teens we spoke with all wanted to be able to experiment with AI tools and learn how to use them effectively and ethically. No one—neither students whose schools allowed ChatGPT and/or other AI tools, nor those whose schools had outlawed them—was in favor of outright bans. But they crave clearer rules around usage, both for learning today in school, and to help prepare them for the future, in which they all agreed AI would play a prominent role. One student from Maryland described a capstone research project for one class where the teacher explained how to properly use AI for research, including appropriate citation methods.

Plagiarism and cheating remain top concerns in schools

One New Jersey 10th grader's school has banned ChatGPT, saying that it's a form of cheating, and warns that teachers have tools that can monitor whether a student has used it. But as the student's mother pointed out, if kids can find a way around it, they will. In other schools, teachers—especially in English and history or social studies classes—are moving to handwritten, paper-based assessments in class to reduce the possibility that kids are using AI. One eighth grader from northern California talked about the elaborate process (all done in the classroom) of writing an essay first in pencil, revising in pen, typing into a Chromebook, printing it, editing the printed version with a pen, entering those edits back into the digital version, printing that final copy, and turning in all of the materials to their teacher in order to document each step of the essay. A seventh grader from Maryland mentioned a classmate who got caught using ChatGPT to write a paper, and had to rewrite it by hand in the classroom.

Students are grappling with ethical concerns, and need more space to learn about AI's risks and harms

Students talked with us about the fear that AI might someday replace humans. They noted that AI could take away jobs from people, but also that the technology might create new pathways of employment. A student from Virginia said, "There's already such a huge focus on STEM, and I feel like the humanities just bring so much more to the table that's often just overlooked, and I feel like this will just kind of further exacerbate that trend, and that's one of my concerns." Bias, inaccuracy, and misinformation didn't come up as topics as much, suggesting that students would benefit from richer discussions in the classroom on these areas. When the New Jersey student's mother asked him if he believes AI is correct, or if he double-checks what he's reading, he said he assumes it's correct. "Wrong answer!" his mom replied.

Social media and chat apps are go-to sources for kids on AI

Information on AI is perceived as easy to access, though students struggled to name specific, reliable sources. Instagram and TikTok often introduce young people to AI via synthetic images. A Maryland 11th grader mentioned a trend around what AI thinks someone from a particular country or particular college looks like. The New Jersey 10th grader finds information on AI via the messaging app Discord, where he has discovered some disturbing AI-generated videos meant to look real. When asked where they'd turn for reliable information on AI, students named large, legacy news outlets such as CNN, the *New York Times*, and the *Wall Street Journal*.

Age matters

Two students felt that age makes a difference in how much generative AI should be used in school. They made the point that younger users need to learn basic skills that AI could circumvent, but older students could take advantage of the technology once those earlier skills have been mastered. On a similar note, the 2024 Impact Research survey found that the older the grades that teachers teach, the more likely that they are to encourage AI chatbot usage.¹⁵⁶

AI is raising existential questions about the definition and purpose of education

A 10th grader in southern California wondered why she's learning computer science if AI will be able to program itself. If AI is the future, what—and how—do students need to learn today in order to be successful in that future world? The Virginia student asked, "What is the purpose of education? Because if it's to learn, then you want to probably minimize AI. But if it's to prepare you for the workforce, then maybe you want to maximize it to improve your efficiency, because that might be what certain employers want in the future."

What is the purpose of education? Because if it's to learn, then you want to probably minimize AI. But if it's to prepare you for the workforce, then maybe you want to maximize it to improve your efficiency, because that might be what certain employers want in the future. — Virginia student

¹⁵⁶ Impact Research. (2024, May). *AI chatbots in schools: Findings from a poll of K-12 teachers, students, parents, and college undergraduates*. <https://www.waltonfamilyfoundation.org/learning/the-value-of-ai-in-todays-classrooms>

Recommendations

We close with recommendations for the field, based on themes and points that emerged throughout our research. These recommendations are intended for a broad audience, which might include educators, students, parents, researchers, industry, policymakers, and funders.

One overarching question that we recommend all stakeholders consider is not only *how* AI technology should be used, but whether it should be used at all. Amid the understandable eagerness to identify promising educational outcomes, this first step often gets overlooked. Is an AI-powered solution, feature, or add-on the right and the best choice in each educational context? As the paper we previously cited on the ethics of AI in education stated, there is a "need to differentiate between doing ethical things and doing things ethically."¹⁵⁷

Bring the doomers and the boomers together, and don't forget the majority in the middle

There's a convener role to play in the AI debate, facilitating meaningful discussion on how to move forward with responsible experimentation while taking critiques into account. We may not reach consensus, but critics should be invited to participate in determining what responsible AI looks like in education. And the vast majority of people who are simultaneously skeptical and curious, trying to navigate this new technology on their own with few guideposts, deserve more attention. They're the real audience, not the extremes with strongly held views.

Listen more carefully to refusal

Beyond critics, an important stance of total refusal exists that we ignore at our peril. Minoritized communities burned by algorithms and other technology in the past may understandably want nothing to do with generative AI; listening to them will only make the conversation more productive and the work ahead more inclusive.

Center equity

We must pay more than lip service to equity, and not assume that bias will automatically get righted as the models evolve. Without intentionality and oversight, it won't. Many of our other recommendations here are tied to equity, such as participatory design and procurement guidelines. As we noted in the section on responsible AI, the values that inform the design of AI systems, platforms, and tools are the same values that inform their outcomes.

We should also not forget historical creative uses of technology among minoritized communities. For example, research shows that Black youth embrace social media as an important means of creative expression more than their White peers. It will be important to acknowledge and elevate ways that overlooked communities in the United States—and communities throughout the global majority—are innovating with AI.¹⁵⁸

Embrace participatory design

The AI table has to expand to include more seats. The idea of participatory design, illustrated in this paper by some of the work at the Giant Room and practiced by many others in the field, can be applied to every facet of AI, including the design and development of technology, the creation of standards and guidelines for schools, the

¹⁵⁷ Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Buckingham Shum, S., Santos, O.C., Rodrigo, M.T., Cukurova, M., Bittencourt, I.L., & Koedinger, K.R. (2022). Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, 32:504–526. <https://doi.org/10.1007/s40593-021-00239-1>

¹⁵⁸ Amrute, S., Singh, R., & Guzmán, R.L. (2022, September 14). *A primer on AI in/from the majority world: An empirical site and a standpoint*. Data & Society Research Institute. <https://datasociety.net/library/a-primer-on-ai-in-from-the-majority-world/>

establishment of an AI literacy curriculum, the building of governance and enforcement mechanisms, and alignment around safety and efficacy metrics.

Note that we use terms such as "design," "creation," and "building." We do so intentionally because we cannot only invite people in to report harm and try to repair it after the fact. The more inclusive and expansive we are at the building stages, the better the outcomes.

We recommend adding so-called seats at the table in three ways. First, by prioritizing marginalized and minoritized communities. We've already seen the ways bias, racism, and classism have infected generative AI. Second, by inviting students, classroom teachers, and families into the process, and not only to give feedback after the fact. Those impacted most by education decisions should be part of the process. The University of Washington College of Education designed and developed a free lesson-planning tool with teachers in the field, for example.¹⁵⁹ And third, the design and development of generative AI tools and platforms in general should be informed by more than technologists. Social scientists, ethicists, human rights advocates, and others can provide important inputs in addition to students, teachers, and parents. When done right, this bigger table expands everyone's frame of reference and provides a more complete set of inputs and perspectives in order to make more ethical, effective, safe, and equitable product decisions.

Invest in AI literacy

All educators need training, and it can't fall to individual teachers, or even individual schools, to seek that out. Districts and states should take the lead to foster a communal sense of learning and understanding and promote shared language and curriculum. Literacy must encompass education about AI and how it works, including its limitations and risks, as well as training on how to use AI tools in and outside of the classroom. Plenty of organizations with terrific resources already exist, including the three we referenced in this paper—the [AI Education Project](#), [AI for Education](#), and [Digital Promise](#)—as well as [Common Sense's own work](#).

Given the proliferation of resources, however, it will be important to map existing offerings according to audience (e.g., students, teachers, administrators, parents), grade level, level of expertise, and focus area (e.g., school policy, creative exploration, bias, misinformation) so that schools can make informed decisions about which training materials best suit their needs at any given juncture.

Standardize guidelines

Schools need clear, specific guidelines on acceptable, appropriate, beneficial, and ethical uses of AI. These guidelines must address academic integrity, procurement criteria, and age and grade level considerations. Teachers need to know what's acceptable for use in the classroom with adult guidance versus what's acceptable for a student to use on their own, and what's OK for them to use as tools to help with lesson planning and curriculum differentiation. Educators and students need to know how to properly cite information found through a generative AI tool. Clear, consistent rules about disciplinary action in the case of inappropriate usage need to be established, with special attention on mitigating the risk of inequitable applications.

¹⁵⁹ Stiffler, L. (2024, May 24). *Chatbots for teachers: Univ. of Washington releases free AI tool for quicker, better lesson plans*. GeekWire. <https://www.geekwire.com/2024/chatbots-for-teachers-univ-of-washington-releases-free-ai-tool-for-quicker-better-lesson-plans/>

Clear guidelines and improved literacy will also help assuage concerns over cheating. The more that teachers know about AI tools and how they work, the better they can determine when a student has crossed a line. Adam Aguilera, the ELA teacher in Washington state, says that teachers themselves are the best AI detectors, because they know their students' writing styles. With proper training, they can detect clues that reveal a swath of an essay, for example, was copied and pasted from a generative AI chatbot.

Adopt procurement standards

School districts and state-level agencies also need to see themselves as customers, wielding real purchasing power. Academic, administrative, operational, and technology leaders need to work together on procurement, not in silos. They can and should demand improvements from tech companies around transparency, data security, privacy, data interoperability, and proof of efficacy. And they should ask vendors about what they're doing to mitigate problems of unfair bias and inaccurate information; even the act of asking sends a signal about the importance of accountability. Control can't only be in the hands of developers and salespeople.

Develop concrete ways to foster transparency and accountability

In addition to legislative and regulatory action, civil society can play a big role in pushing the tech industry on transparency and accountability. Common Sense Media's own ratings and reviews are a start; we need more "nutrition label" information formats for educators, students, and families to be able to make decisions. Digital Promise has also launched certification processes for companies in research-based design, learner variability, practitioner-informed design, and learning and employment record technologies.¹⁶⁰ They're also working with University College London on a trust certification, which would show that a company is oriented toward efficacy and impact.

Consider how generative AI may impact both learning goals and education's higher purpose

Educators shouldn't start with a shiny new AI product and then figure out how it fits into the classroom. As we saw with professor Jason Yip's example from his own class in informatics, teachers need to focus on the goal of a lesson or curriculum, and then determine how generative AI might support or disrupt that goal. Stepping back, we also need more conversation around the learning ecosystem writ large—to understand how generative AI impacts how students learn, what they're learning, and why they're learning what we're teaching them. As the 11th-grade member of Common Sense's Youth Research Council asked, "What is the purpose of education?" Generative AI's emergence should make us pause and reexamine that question.

Exercise caution—and practice the lost art of patience

Remember that generative AI is still experimental technology, and that the public has essentially served as test subjects in its development. Given how much we don't yet know about the long-term effects of any generative AI tools, and the harmful effects of the "move fast and break things" ethos that we do know about, adoption and use in education should proceed carefully. Adequately preparing young people for the future shouldn't come at the expense of exposing them to harms today.

Conclusion

When we set out on our landscape analysis, we knew there was a lot to cover—and that we had to do it quickly. The field of AI, and generative AI in particular, is rapidly evolving and expanding, with new platforms, products,

¹⁶⁰ Digital Promise. *Product certifications*. <https://productcertifications.digitalpromise.org/>

updates, frameworks, and policy proposals popping up faster than any human can keep up with (perhaps another use case for an AI tool). We've therefore tried to focus on the biggest themes that we know won't disappear in the coming months.

We hope our analysis illuminates this landscape and clarifies the stakes. By providing this survey of definitional material, public opinion, policy overviews, opportunity highlights, risk assessments, and recommendations, we also hope that readers come away better equipped to make what will no doubt be pivotal decisions around AI in education. And by including the voices of different stakeholders—students, parents, educators, researchers, and product developers—we hope we've underscored the most important theme of all: humanity.

Acknowledgments

Common Sense Media is grateful for the support of the William and Flora Hewlett Foundation that enabled this work.

Report research and writing: Bene Cipolla, with support from Amanda Lenhart

Copy Editors: Christopher Dare and Jennifer Robb

Designers: Chris Arth and Emely Garcia

We would like to express our gratitude to the following people, all of whom were very generous with their time and expertise: David Ball, Amanda Bickerstaff, Jean-Claude Brizard, Dan Carroll, Ben Caulfield, Roderic Crooks, Dora Demszky, Jimmy Fisher, Jennifer Frentress, Laurence Holt, Azadeh Jamalian, Alex Kotran, Alex Mahadevan, Rajiv Maheswaran, Chantell Manahan, Babak Mostaghimi, Erin Mote, Zachary Pardos, Kelsey Peterson, Rajen Sheth, Emily Tavoulareas, and Jason Yip.

We'd also like to thank the many young people, parents, and teachers who shared their lived experiences for this paper.

From Common Sense, we would also like to thank Amy Brotman, Eisha Buch, Daniel Vargas Campos, Sara Egozi, Amina Fazlullah, Tracy Pizzo Frey, Holly Grosshans, Jeff Graham, Ariel Fox Johnson, Girard Kelly, Ben Kornell, Marisa Naughton, Ellen Pack, Sarah Rosenbach, Beth Sears, Ashwin Sridhar, and Robbie Torney for their insights and contributions.

Finally, we are grateful to Adam Aguilera, English Language Arts Teacher, NBCT; Janet Haven, Executive Director, Data & Society; and Pati Ruiz, EdD, Senior Director of EdTech and Emerging Technologies at Digital Promise, for reviewing this report.

Interviews for this paper were conducted from December 2023 to March 2024, with some follow-up conversations and external review in April 2024.

Appendix

Resources, links, and information

Common Sense product ratings and lists

[AI Initiative | Common Sense Media](#)

[How we review and rate AI products | Common Sense Media](#)

[Classroom Tools That Use AI | Common Sense Education](#)

AI literacy

[AI Literacy Curriculum Hub](#): a roundup of various education-focused resources from AI for Equity and [Erica Murphy](#) of [Hendy Avenue Consulting](#)

AI literacy lessons for students

[AI Literacy Lessons for Grades 6–12 | Common Sense Education](#)

[Teach AI – aiEDU](#)

[Curriculum – AI for Education](#)

[Activity Resource Guides for Teaching Artificial Intelligence in K–12 – AI4K12](#)

AI literacy trainings for educators

[Professional Learning – aiEDU](#)

[Introduction to Gen AI for Educators PD Workshop](#)

AI literacy framework for families

[The 4As: Ask, Adapt, Author, Analyze – AI Literacy Framework for Families](#)

[Administrative Foundational Questions for Using Generative AI Tools](#)

Educator snapshot: how one social studies teacher is navigating AI

This interview with a high school social studies teacher in Huntington, New York, who asked to remain anonymous for privacy reasons, offers a glimpse into one educator's experiences with AI. We include it to provide human context and the perspective of someone working in the classroom and trying to navigate new technology.

Q: How familiar are you with the term "artificial intelligence"? What does it mean to you?

Teacher: I am comfortable with technology in general, but I wouldn't call myself an AI expert by any means. We've had one PD [professional development] session regarding it. AI, to me, I think of it as generative AI—ChatGPT and that sort of thing.

Q: What was the PD session like?

Teacher: It was very introductory, very eye-opening as well. I actually hadn't dabbled with ChatGPT until that point. It was eye-opening to see how powerful it can be. It reminded me of the old saying, "With great power comes great responsibility." It was incredible to me that you could request to the GPT bot, "Create a 10-question, four-choice multiple choice quiz on the topic of whatever," and it came right up. Not all the questions were great, but there were some decent ones, and you could use those to springboard forward.

Q: What kinds of technology do you use in your classroom? Do you use any AI tools?

Teacher: I learned about MagicSchool in the PD session, which helps teachers. I've dabbled with MagicSchool to create a worksheet. I have encouraged students to use ChatGPT to help quiz them ahead of tests.

Q: What about your colleagues?

Teacher: It's not widely discussed. Just because it's new and people are aware of concerns of plagiarism.

Q: Are you familiar with any of the risks with this technology?

Teacher: There's chatter about *Terminator* [movie reference] and that sort of thing. We just had one PD session, so there wasn't a lot of opportunity to explore in depth. I imagine more PD is coming. That was just an introduction this year.

Q: Does your school have guidelines on AI usage, including anything on ChatGPT in particular?

Teacher: For teachers, none that I am aware of. For students, there are concerns with plagiarism, but I don't know if they have been fully addressed. We have tools for sussing out traditional plagiarism—copying from known sources or copying from student papers. With generative AI, there's nothing to check against it. It's very delicate. You might recognize it's not their voice—different word usage, for example—but I don't know of anybody that's gone forward with a full-throated accusation of generative AI plagiarism. It's very difficult to prove, and I'd be concerned about pushback from parents without some definitive proof.

Q: What are your concerns with generative AI?

Teacher: We've seen this evolution in technology that has always been a tool to make us more productive, whether going back to the adoption of calculators in schools or turn-by-turn directions in cars. Super helpful. But people *have* become less adept at doing basic calculations or directions. The question in my head: Is this going to help us be more productive? I'm worried about technology replacing human endeavors—that's super scary. For a long time in education, we've been talking about teaching kids today so they can be prepared for jobs that don't

exist today. But what if jobs writ large don't exist? And I don't only mean unskilled labor. Accountants, lawyers, teachers—these are going to change radically in the future, to say the very least.

Q: Anything else you'd like to add?

Teacher: When there's this new technology, people have wild ideas about what may come of it. I remember reading a book called *The Devil in the White City*, and one of the things that blew my mind was that their understanding of technology at the time had them predicting absolutely wild things that would exist in the future—things that we laugh at now. I sometimes wonder whether we are in a similar phase with this. There's a lot of talk and discussion, and people have expressed serious concerns about the future. I don't want to be Chicken Little, but I also don't want to be caught unawares.

Educator snapshot: how one fourth grade teacher uses Kyron Learning

Lizzy Brent teaches fourth grade at a public charter school in California. The school was one of 46 piloting Kyron Learning during the 2023–24 school year, which was Brent's first as a teacher. We include this interview to provide a snapshot of how teachers are using, navigating, and reflecting on AI.

Q: How does the tool work?

Lizzy Brent: The students log in, they get to choose a teacher. I assign a lesson on Kyron based on a skill I believe needs to be targeted. I use it like a reteach when I'm working with a small group of students, and also for looking ahead. If we're going to be doing division or fractions, I'll go through Kyron and pull a lesson so they have more background knowledge when I actually teach it.

What I like about Kyron is that they gave us a document where they show how their lessons align [with different curricula]. We use Bridges [math curriculum]. Kyron can teach the same concept, but in a different way. So it provides, I think, different opportunities for student processing.

Q: How do you introduce the lesson?

Brent: I post it on Google Classroom. I'll say, "OK, students, when you go to Google Classroom under Math, you'll see a Kyron lesson." They need to have a whiteboard with them so they can show their work and do the work on their whiteboard while they're following along with the virtual teacher.

Q: Do the students show the whiteboard to the camera?

Brent: No, so the way that it works, it's almost like a prerecorded lesson with the teacher. However, the teacher will pause and ask questions, like "What do you get for this?" or something. And the students have to type in their answer. Based on the student answer, if they get it correct, it'll move forward. If they get it incorrect or there's a misconception, it'll reteach until it can find where the student misconception stems from.

So they have to type in their answers, but what I've noticed is when students are on a computer and they're inputting answers, it's not automatic for them to use scratch paper. They think they have to do it in their head and then type in an answer as quick as possible. So I have them show their work on their whiteboard for accuracy, so they're actually applying the strategies that are being used.

Q: Do you use Kyron for homework or just in the classroom?

Brent: Only in the classroom. Technically, I could assign it as homework because students have access to Google Classroom at home, but I'm still very new to Kyron, so I'm just trying to get to mastering it in the classroom.

Q: What have you noticed with the tool? What are you seeing?

Brent: One thing that I'm liking is student independence. And the students get really excited about choosing their own teacher, and I've had some students choose to learn the concept in their home language [Spanish] instead of in English. All of my students are multilingual learners.

Q: Anything else you're seeing in terms of learning?

Brent: I don't have hard data about before we used Kyron, and there's so many factors for every student, right? But on the teacher interface, I can see the lessons that I've assigned and I can see how much students have progressed and how much they've fallen behind. I mean, it literally shows where they are in the lesson. And I just noticed, it's kind of like the same thing, the students that usually work slower, I see them still working slowly through Kyron. The students that are typically faster are getting to the end faster. But it also shows me all the misconceptions. So if I look at this lesson that I assigned recently, about using a number line to add two fractions with the same denominator, and I look at the misconceptions tab and I see that three students had this misconception, four students had this misconception, and it groups them for me and it has the student names. And so I can see that 18 of my students had a misconception on labeling the number line.

Q: How do you use that information?

Brent: That's something that I'll need to go back and reteach. I think, in that way, it supports the students. I'm assuming that it's helping the students because it's helping me make more clear, refined decisions. I also have to look at their answers and figure out, was it an error because they didn't input the answer correctly on the computer, or is it because they didn't really understand the concept? I noticed with a couple of my multilingual learners that were doing it in English, they weren't using their home language, when they were in the lesson, [and] sometimes they didn't understand the question.

Q: Do you use any lesson planning tools, or anything that's directly supporting a teacher?

Brent: Two days ago I heard about MagicSchool. I've used it once and I definitely want to go back and use it. But the thing about it is, you have to be specific in the words you use because sometimes the questions that it generated, I was like, "Ooh, that is not a fourth grade question."

Q: So you still have to double-check everything.

Brent: Absolutely, you have to triple-check everything.

Q: What about at home, have you used ChatGPT or Gemini or any other tools?

Brent: I know other people have, there's some people I know that live on ChatGPT, but I'm good with writing my own essays. So I don't use a lot of AI.

Q: Have you gotten any sort of training or PD via the school on AI?

Brent: No.

Q: Do you feel literate on AI, how it works or the risks of the technology?

Brent: I definitely have my own understanding and I have my own opinion about AI. I'm actually really concerned about it. I think it's great for certain things, but I'm concerned that it's taking away the productive struggle for people. If you're just trying to find shortcuts using a computer to do the thinking for you, I think that's doing a disservice and it's not supporting our need to learn to critically analyze and stuff.

Educator snapshot: how one fourth grade math teacher uses TeachFX

Raven Johnson teaches fourth grade math at L.P. Miles Elementary School in Atlanta. As grade chair, she was part of a small cohort testing TeachFX during the 2023–24 school year. TeachFX provides AI-powered instructional feedback. We include this interview to provide a snapshot of how teachers are using, navigating, and reflecting on AI.

Q: How does TeachFX work? How do you use it?

Raven Johnson: You can use it on your iPad or phone. There's a better connection on the phone. It's an app that allows you to record lessons. It monitors how much you and your students talk, and how much group time talk you have. Also dead talk. It takes everything you and your students say. Then it gives you an overview of the lesson. It tells me what I taught, it gives me a visual representation that I can use with the kids to do a reflection about what we learned. It gives me feedback and questions, like ideas of what I can ask.

Q: How does it help you teach?

Johnson, pointing to the app on her phone screen: This makes me feel good because I learned I asked 66 questions. It helps track data to see how much you have your kids talking.

Q: Do you have a student talk time target?

Johnson: I teach 84 kids a day. I don't have a specific student talk time goal. Next year I might set a goal.

Q: What have you learned about your teaching? Have you made any changes as a result of using this tool?

Johnson: I used to get feedback about letting kids talk more. When I first started using TeachFX, it was showing me I was talking a lot. Since I started monitoring, I let the kids do the work more, let the kids come to the board more instead of doing the work for them.

Q: Do you share the results of the tool with your principal or a coach?

Johnson: No, unless they ask. This is our first year, so they kind of let us venture off this year on our own. They haven't really asked for it, so I just use it for my own benefit. I may catch an *a-ha* moment and download it to my phone, like I remember one time a student was like, "This is so fun and thank you." Now I can hear that forever.

Q: What other digital tools do you use in the classroom, if any?

Johnson: ClassDojo to communicate with parents. Kids can earn points to get a snack on Thursdays or a bigger reward at the end of the month. We use Google Classroom. If we play a math game, we use a digital spinner.

Q: How familiar are you with how these digital tools work? What's your comfort level with tech in general, and with generative AI specifically?

Johnson: I play around with ChatGPT on my own time. Even when I am writing an email, I'll run it through GPT to help me make it better. I make flyers on Canva, lessons and worksheets. I've tried to play around with teacher AI tools but they take all day, so I just do it myself. There's so many out there.

Q: Do you talk about AI with your colleagues? Has your school or district provided any guidelines on the use of AI?

Johnson: Colleagues have shared some websites with me, especially science teachers. They haven't said anything about their policies for using it. I don't think the kids at this young age know about it.

Q: Do you have concerns about AI?

Johnson: I'll be 26 this month. We had to work hard for a lot of stuff in school. Now that AI is taking over, how will they teach this new generation to be independent? Will they know how to do things without AI? And how would they stop cheating when writing papers, plagiarism, things like that?

About Common Sense Media

Common Sense is the nation's leading nonprofit organization dedicated to improving the lives of all kids and families by providing the trustworthy information, education, and independent voice they need to thrive in the 21st century. Our independent research is designed to provide parents and caregivers, educators, health organizations, and policymakers with reliable, independent data on children's use of media and technology and the impact it has on their physical, emotional, social, and intellectual development. For more information, visit commonsense.org/research.