

ARTICLE

State Education Policy and the New Artificial Intelligence

The Technology Is New, but the Challenges Are Familiar



AUTHORS

Glenn M. Kleiman | Stanford Accelerator for Learning

[H. Alix Gallagher](#) | Policy Analysis for California Education, Stanford University

PUBLISHED: September 19, 2024

NASBE | National Association of State Boards of Education

This article is intended to provide State Boards of Education nationwide with information about [artificial intelligence \(AI\)](#) and how it can be used in education. Understanding generative AI's strengths and limitations is crucial for making informed decisions about how to integrate these tools effectively while maintaining the integrity of human-centered education,

[Generative AI](#) differs from traditional computing because it learns from and creates based on vast datasets, using machine-learning algorithms to form neural networks rather than following step-by-step programming. Generative AI can therefore produce novel responses autonomously, often in ways that are unpredictable even to its creators. Optimists believe AI will enhance the role of teachers by providing customized learning resources, digital tutors, and innovative experiences tailored to individual students' needs. In this vision, AI works alongside teachers, enabling them to focus on building relationships, understanding diverse student backgrounds, and inspiring learners. This partnership—where humans remain at the center of teaching and learning while AI assists in the process—could help close opportunity and achievement gaps, promoting meaningful educational reforms that prepare students for a rapidly evolving world. Conversely, pessimists express concerns about the potential dehumanization of education. They argue that AI could increase students' reliance on digital tools, reduce meaningful human interactions, and perpetuate biases and misinformation. Critics also fear that AI could widen equity gaps, compromise student privacy, and encourage disengagement by enabling students to misuse AI for academic tasks.

The article highlights the need for education leaders and policymakers to navigate the use of AI with care, balancing its transformative potential with its inherent risks.

Introduction

In 1953, a [UNESCO](#) report titled *Television: A Challenge and a Chance for Education* offered the following views about the use of the nascent technology in schools:

- It opens a new era of educational opportunity.... Something new has been added. That something new is exciting, it is challenging, it is revolutionary.
- It is a social force with enormous potentialities for good or bad.
- It is, at the present time, probably overrated by pioneer enthusiasts, as it is also too summarily dismissed by others.
- It should be used to complement rather than displace the teacher.
- The condition of further effective advances in this field is caution—great deliberation, more careful and extensive testing and, above all, concern with standards.

Similar statements arose in past debates about calculators, personal computers, the internet, mobile devices, and virtual classes. They echo again in the abundant optimistic and pessimistic views about the potential educational impacts of [artificial intelligence](#) (AI) tools.

The optimists envision that AI will enable teachers to do more of what only teachers can do for their students: build caring and trusting relationships; understand students' needs, backgrounds, and cultures; guide and inspire their learning; and prepare them for their futures. In this vision, AI partners with teachers to provide customized learning resources, digital tutors, and new learning experiences while addressing students' individual needs. It helps address long-standing opportunity and achievement gaps among student groups. Overall, it sparks reforms that go beyond increasing the efficiency of current practices, enabling significant improvements in the quality of teaching and learning so all students are prepared to succeed in the rapidly changing world.

The pessimists envision that students using AI will spend more time interacting with digital agents than with human teachers and peers. In this view, AI educational resources are rife with biases, misinformation, oversimplification, and formulaic, boring presentations. It sees students misusing AI to do their work for them, resulting in a lack of engagement and productive learning. It sees AI enlarging equity gaps and jeopardizing students' privacy and security. It sees policymakers replacing teachers with AI to address funding shortfalls. Overall, the fearful vision is that AI will lead to the dehumanization of education.

Education leaders and policymakers will need to consider the differing views in order to harness the potential and mitigate the hazards presented by the advances in AI technologies.

What Is Generative AI and What Can It Do?

To make informed decisions about generative AI, one needs to understand how it differs from prior technologies, its limitations and risks, how it relates to human intelligence, and what it can do now and will do in the foreseeable future.

How is *generative AI* different from prior technologies? Before the advent of generative AI, computers had long been described as instruction-following devices that required detailed, step-by-step programs. A computer, we all learned, could only do what it was programmed to do, and errors occurred if there was a flaw—a bug—in the program. How a programmed computer responded to inputs to produce its outputs was well understood, and one could trace how the inputs were processed through the steps of the program.

In contrast, generative AI systems learn and create. They are built with machine-learning algorithms that extract patterns from enormous data sets. The extracted information is captured in digital neural networks that are somewhat analogous to the neurons and synapses of the human brain but differ from human brains in many ways. The scale of these systems is difficult to fathom. The texts used to train recent [AI large language models](#) (LLMs) are equivalent to millions of books, and the resulting neural network has hundreds of billions of connections—fewer than, but beginning to approach, the number in the human brain.

Once the neural network is created, requests to the AI model trigger algorithms that use the network's information to generate responses. That is, a generative AI system builds a knowledge structure—the neural network—from the data on which it is trained and then uses that knowledge to respond to a wide range of requests. How generative AI produces its outputs is not at all transparent, any more so than the human brain. AI models can process inputs and generate outputs without being programmed with step-by-step instructions, creating novel responses each time. Since generative AI models are not limited to tasks for which they have been programmed, they often do things that surprise even their creators and other technology leaders. For example, Bill Gates reported that he was "[in a state of shock](#)" when he first realized what these systems can do.

What are the limitations and risks for education? While generative AI can do many things, it has limitations and risks that are especially concerning when used in education:

- LLMs generate responses from patterns in their training data, but they do not verify the accuracy of that data. Thus AI models can fabricate information, often convincingly, resulting in what are called hallucinations.
- Websites and social media included in LLM training data sets may contain erroneous, biased, conspiratorial, violent, and other inappropriate and toxic information that can be incorporated into neural networks and appear in AI-generated outputs.
- AI models can respond inappropriately to the social and emotional states of users, so they can provide damaging responses to requests about significant problems such as depression and suicidal ideation.
- AI models can be used intentionally to deceive and misinform by enabling people to create realistic, deepfake pictures, audio recordings, and videos that make it appear that someone said or did things they would never do.
- When students use generative AI, it raises concerns about privacy and security and about violating the federal [Children's Online Privacy Protection Act](#) (COPPA), [Children's Internet Protection Act](#) (CIPA), and [Family Educational Rights and Privacy Act](#) (FERPA) requirements.

The heart of the matter is that generative AI does not distinguish constructive, accurate, appropriate outputs from destructive, misleading, inappropriate ones. AI developers and researchers are actively working to mitigate these problems. They are seeking to better curate the training data, provide human feedback to train the AI models further, add filters to reduce certain types of outputs, and create closed systems that limit the information in a neural network. Even so, the problems will not be entirely eliminated, so educators need to be aware of them and plan to protect students and teachers from potential harms.

How does generative AI relate to human intelligence? With its conversational chat style, generative AI mimics human interactions, so it is easy to anthropomorphize it. But understanding how artificial and human intelligence differ is essential for determining how AI and people can best work together.

First, each learns differently. LLMs can be said to have only book learning—all they "know" is derived from the enormous amount of text on which they are trained. In addition to learning from text and other media, human learning involves physical explorations of the world, a wide range of experiences, goal motivation, modeling other people, and interactions within families, communities, and cultures. Human infants have innate abilities to perceive the world, interact socially, and master spoken languages.

Human learning is ongoing. It involves understanding of causes and effects, human emotions, a sense of self, empathy for others, moral principles, the ability to understand the many subtleties of human interactions. In contrast, LLMs generate outputs through a mathematical process involving the probability of what might come next in a sequence of words: There is no intuition or reflection. AI can mimic human outputs, but the underlying processes are very different.

Star Trek's Captain Kirk and the half-alien Mr. Spock in the original series—and Captain Picard and the android Data in the later series—provide a helpful analogy, as suggested by education researcher [Chris Dede](#). The human captains bring judgment, decision making, instinct, deliberation, passion, and ethical concerns, while the logical, nonemotional Spock and Data bring analytic and computation skills. Working together enables the nonhuman intelligence to augment the powers of humans so they can overcome

the many challenges they face while traveling through space.



Photo credit: [Wikimedia Commons](#).

This *Star Trek* example illustrates the concept of intelligence augmentation, which the U.S. Department of Education emphasizes in its recent report on AI, teaching, and learning. In this approach, humans lead. They set goals and directions, and AI augments their capabilities so that humans and machines together are far more effective than either alone. This perspective recognizes the unique roles of teachers and keeps them at the center of student learning.

What can generative AI do? Advances in generative AI triggered an avalanche of news reports about its impact on human work, learning, communication, creativity, research, decision making, and civic life. LLMs' ability to write messages, blogs, reports, poems, songs, plays, and other texts have received the most attention. AI systems and tools can also produce works of art, analyze X-ray and MRI images to help diagnose medical problems, make predictions and recommendations to inform business decisions, solve scientific problems such as predicting the structures of proteins, drive cars through busy streets, enable knowledge workers to be more productive, and more. AI is already changing many industries and occupations, and analysts have projected more changes are in store.

Generative AI tools can play many roles in schools. Teachers are already using them to help create personalized lesson plans. Administrators are communicating with their communities using AI's ability to draft messages and translate across languages. Students are working with AI tutors, using AI in research and writing projects, receiving AI feedback on draft essays, and learning from AI-enabled simulations.

Wharton School researchers [Ethan Mollick](#) and [Lilach Mollick](#) describe seven roles that AI can play with students: mentor, tutor, coach, teammate, student, simulator, tool. They summarize the benefits and risks of each (table 1).

Table 1. AI Roles in Education

AI Use	Role	Pedagogical Benefit	Pedagogical Risk
Mentor	<i>Providing feedback</i>	Frequent feedback improves learning outcomes, even if all advice is not taken.	Not critically examining feedback, which may contain errors.
Tutor	<i>Direct instruction</i>	Personalized direct instruction is very effective.	Uneven knowledge base of AI. Serious confabulation risks.
Coach	<i>Prompt metacognition</i>	Opportunities for reflection and regulation, which improve learning outcomes.	Tone or style of coaching may not match student. Risks of incorrect advice.
Teammate	<i>Increase team performance</i>	Provide alternate viewpoints, help learning teams function better.	Confabulation and errors. "Personality" conflicts with other team members.
Student	<i>Receive explanations</i>	Teaching others is a powerful learning technique.	Confabulation and argumentation may derail the benefits of teaching.
Simulator	<i>Deliberate practice</i>	Practicing and applying knowledge aids transfer.	Inappropriate fidelity.
Tool	<i>Accomplish tasks</i>	Helps students accomplish more within the same time frame.	Outsourcing thinking, rather than work.

Source: Ethan R. Mollick and Lilach Mollick, "Assigning AI: Seven Approaches for Students, with Prompts," Wharton School Research Paper, forthcoming (SSRN, September 23, 2023).

AI continues to advance rapidly. For example, it is now being integrated into word processors, spreadsheet programs, graphics programs, and instructional resources. AI agents are becoming common: They can be assigned a goal, develop a plan to fulfill the goal, take actions to achieve the goal, and then monitor progress and adjust the plan as necessary.

Not only does generative AI differ from prior technologies in how it works and what it can do, it differs in how people can interact with it, as is eloquently described by scholar [Punya Mishra](#):

// ... GenAI doesn't just operate in isolation, but it interacts, learns, and grows through dialogue with humans. This collaborative dance of information exchange collapses the old boundaries that once defined our relationship with tools and technology... [H]ow we make sense of these new tools is emergent based on multiple rounds of dialogue and interactions with them. Thus, we're not just users or operators, we're co-creators, shaping and being shaped by these technologies in a continuous and dynamic process of co-constitution. This is a critical shift in understanding that educators need to embrace as we navigate the wicked problem of technology integration in teaching.

The challenge for state education leaders is supporting the thoughtful exploration and adoption of tools that improve teaching, learning, and school and district functioning, while AI innovation continues to advance at a rapid pace.

Recommendations for State Education Leaders

AI will impact all aspects of schooling, from curriculum to assessment, from student services to teaching practices, and from school management to community engagement. State leaders will have many decisions to make about policies and programs to foster productive, appropriate, safe uses of AI. Here are some general recommendations to help education leaders consider how to proceed.

Focus on how AI can help educators address critical challenges. The rapid advances in AI have led educators and policymakers to ask, What should we do about AI? We recommend asking instead, What are our critical educational challenges, and can AI help us address them? For example, how can AI help in improving teacher working conditions in order to attract and retain teachers, increasing students' engagement in learning and achievement, reducing long-standing opportunity and achievement gaps that the pandemic exacerbated, and enabling administrators to be more effective leaders, managers, and communicators?

Respond to immediate needs while also looking to the future. Some decisions require immediate action, such as addressing safety and privacy requirements for AI tools in education. Others involve establishing ongoing programs, such as professional learning to prepare teachers and other staff to use AI productively. Others are longer term, such as updating curriculum goals, pedagogical approaches, and student assessments to incorporate AI. AI is not a short-term problem to be solved; it requires processes for ongoing analysis, evaluation, and decision making.

Build upon existing policies and programs. We recommend reviewing and updating existing policies and programs to address AI rather than creating AI-centric ones. Many of the concerns and needs AI is driving resemble what is already addressed in existing policies and processes for the acceptable use of technology, academic integrity, professional learning, digital security and privacy, teacher evaluation, procurement of technology and instructional materials, and other areas.

Make use of resources designed to guide educational leaders. Resources are already available to guide state leaders, and more are being prepared. [TeachAI](#) is a collaboration of many nonprofits, government agencies, technology and education companies, and experts (including ourselves) who work together to provide such resources (figure 1). Their policy resources address five needs:

- **Foster** Establish an [AI in Education Task](#) Force to oversee policy development and implementation.
- **Provide** Equip schools with guidance on the safe and responsible use of AI. States' education departments have been [releasing guidance](#) for their schools, with 15 having provided initial guidance as of June 2024.
- **Build** Provide funding and programs to support educator and staff professional development on AI.
- **Promote AI** Integrate AI skills and concepts—including foundational principles, social impacts, and ethical concerns—into existing instruction.
- **Support** Promote the research and development of safe, effective AI in education practices, curricula, and tools.

TeachAI also produced an [AI Guidance for Schools Toolkit](#) that can help guide the support the state provides to schools and districts. The toolkit addresses the benefits and risks of AI (figure 1).

Figure 1. Potential Benefits and Risks of AI in Education

Potential Benefits				
Content development and differentiation	Assessment design and timely, effective feedback	Tutoring and personalized learning assistance	Aiding creativity, collaboration, and skill development	Operational and administrative efficiency
Potential Risks				
Plagiarism and academic dishonesty	Diminished student and teacher agency and accountability	Compromised student privacy and unauthorized data collection	Overreliance and loss of critical thinking	Perpetuating societal bias

Source: TeachAI, [Guidance for Schools Toolkit](#), 2023.

Support learning about computer science and AI. State boards of education are typically responsible for teacher credential requirements, programs to upskill existing teachers, and the standards and frameworks that specify what students should be learning. Adapting existing requirements for educators and students will catalyze shifts necessary to prepare students for a future where AI has a prominent role. Additionally, the guidance and broader documentation that will be created through these processes can themselves be tools for the state board to communicate key ideas broadly to all levels of the educational system and to the families served.

Keep calm, engage stakeholders, and plan carefully. Corporate, government, and university research centers worldwide are making enormous investments and efforts to advance AI. New AI products become available every month.

The remarkable advances and overwhelming attention to generative AI technologies make it likely that we will witness the hype cycle that has occurred with prior technologies, which has been well documented and described by the [Gartner](#) consulting firm. In this cycle, a technological advance can first *trigger inflated expectations*. In education, this leads to rediscovering that simply introducing a technology into schools is insufficient to producing the desired results. The cycle then enters the *trough of disillusionment*. But some innovators continue to see possibilities and work to instantiate them, providing models that influence others and placing the use of the technology on the gentle upward *slope of enlightenment*. As the effective use of the technology becomes more widespread, it moves to the *plateau of productivity*, in which it becomes embedded throughout classrooms, schools, and districts. Reaching this stage in AI use will require that providers and educators mitigate the risks involved with current AI systems.

In contrast to AI’s rapid pace of change, educational change requires years of thoughtful work to reform curriculum, pedagogy, and assessment and to prepare the education workforce to implement the changes. The difference in these paces leaves educators and policymakers feeling they are behind and need to catch up. We recommend a “keep calm and plan carefully” mindset to foster the careful consideration, thoughtful implementation, and involvement of multiple stakeholder groups that are required for changes in education policy to successfully harness the benefits of AI and minimize the risks.

Authors [Glenn M. Kleiman](#), Stanford University [Accelerator for Learning](#), and [H. Alix Gallagher](#), Policy Analysis for California Education, Stanford University, thank their colleagues Barbara Treacy, Chris Mah, and Ben Klieger for thoughtful comments on a prior draft of this article. This article is reposted by PACE courtesy of State Education Standard: The Journal of the National

Association of State Boards of Education, where it was *published* in volume 24, issue 3. The text appearing on this page is verbatim from the initial publication with the addition of relevant embedded links and follows the content and style standards of NASBE publications.



Stanford Graduate School of Education

520 Galvez Mall, Suite 444

Stanford, CA 94305

Phone: 650.576.8484

edpolicyinca.org

