

Learning Systems

Opportunities and Challenges of Artificial Intelligence-Enhanced Education

By Amy Chen Kulesa, Michelle Croft, Marisa Mission, Brian Robinson,
Linea Harding, Mary K. Wells, Andrew J. Rotherham, and John Bailey

SEPTEMBER 2024 *(Updated August 2026)*

Note: *This report reflects updates through June 15, 2026.*





Contents

Introduction

Key Takeaways

Overview

Platform

Policy

User Experience

Conclusion

Sources

Acknowledgments

About the Authors

About Bellwether

Introduction

Generative artificial intelligence (GenAI) advancements have prompted both excitement and concern about how AI could be used in education. From streamlining administrative tasks to personalized resources for students with learning differences to reimagining staffing models, AI offers a spectrum of possibilities. However, its integration requires careful consideration of both potential benefits and inherent risks.

This three-part series, developed through consultations with experts and practitioners, aims to take a holistic approach, analyzing the opportunities, risks, and considerations surrounding AI in education in the following ways:

Foundations

The Landscape of Artificial Intelligence in K-12 Education

An introductory guide exploring AI developments, capabilities, potential impacts, and current usage.

Considerations

Opportunities and Challenges of Artificial Intelligence-Enhanced Education

A deep dive into the broader ecosystem, including technology platforms, policy landscapes, and stakeholder experiences.

Applications

Artificial Intelligence Use Cases

An outline of how AI could potentially be used in education, detailing opportunities and risks for each use case.



Contents

Introduction

Key Takeaways

Overview

Platform

Policy

User Experience

Conclusion

Sources

Acknowledgments

About the Authors

About Bellwether

Key Takeaways

Insights From the Learning Systems Series

The education sector has an opportunity to shape its future as GenAI technology is developing. Proactive planning allows thoughtful integration of AI’s potential benefits to carefully address its challenges and prepare for emerging capabilities. To help education stakeholders navigate this complex and evolving landscape, Bellwether offers the following recommendations to build strong capacity, resilient infrastructure, and thoughtful design.

Building Strong Capacity

To safely and effectively incorporate AI, the education sector must invest in the development of knowledge, skills, and diversity of leaders and contributors. **Capacity-building effort is not just beneficial — it is critical.** While these efforts require a meaningful investment of time, talent, and resources, without them, the education sector faces barriers to progress. Even though capacity-building around AI may not rise to high urgency for individual organizations, funders and policymakers can and should incentivize the system so that there is meaningful progress over time.

1. Strengthening AI literacy.	Administrators, teachers, students, and families alike need to build understanding of AI’s capabilities, limitations, and implications so that communities can make informed decisions about AI and effectively choose where to (and not to) integrate AI into learning environments. AI literacy also extends into the thoughtful teaching of AI — what students should know and be able to do related to being prepared for the workforce of the 2030s.
2. Enabling the development of policies and guidelines.	Comprehensive and flexible policies and guidelines are essential to wider AI adoption but have been slow to develop, notably at the federal and state levels. As of August 2024, approximately half of states have yet to provide AI guidance. State and local education leaders would benefit from collaboration opportunities, including with industry leaders, as they develop and refine their policies.

Building Strong Capacity *(continued)*

3. Bridging educators, developers, and researchers.	Currently, there is a gap between teams building AI tools and education practitioners. Fostering strong partnerships, especially in development and design stages, can ensure AI tools are safe, effective, and relevant for education purposes, aligned with pedagogical principles, and grounded in research about how students learn. Partnerships will increase the likelihood of well-designed solutions, thereby enabling the realization of AI's transformative potential in education.
4. Expanding sectorwide capabilities.	The rapid advancement of AI in education has surfaced a significant gap between current capabilities and needed expertise. To bridge this chasm, the sector needs to expand capacity across the entire educational ecosystem, addressing needs such as data privacy, ethical implementation, procurement, and other challenges surfaced throughout this three-part series. This expansion can take multiple forms: Existing education organizations can develop AI competencies, AI experts from other sectors can be encouraged to apply their knowledge to educational contexts, and new organizations can be established to address emerging needs. Addressing these gaps in the education sector is crucial for developing comprehensive, education-specific AI solutions that incremental improvements in existing structures may not be able to achieve.
5. Diversifying voices at the table.	Degree attainment data suggest that the current AI field is not yet reflective of the diverse students and educators across the country. Diversification of voices involved in AI-related dialogues and decisions can promote more inclusive AI solutions that benefit all students. This includes conversations with and among AI developers, solution builders, and researchers as well as authentic engagement with diverse educators, students, and families who will be most impacted by decision-making processes.

Building Resilient Infrastructure

In the excitement of thinking about how AI can be used, there is the potential to overlook the importance of a robust, secure, and nimble infrastructure. **Strong data and research not only enhance the quality of AI outputs but also serve as a crucial safeguard, particularly in protecting sensitive student data.** Investing in resilient infrastructure is essential for realizing AI's benefits while mitigating its risks, ensuring that enthusiasm for innovation is matched by a commitment to security and quality.

6. Robust data infrastructure.	Implementing AI widely throughout a system or organization will require a secure, high-quality data infrastructure. This includes ensuring data quality and usability, system interoperability, robust security measures, well-planned change management processes, and reliable internet connectivity.
7. More quality datasets and new benchmarks.	To enhance the quality of AI outputs, there is an opportunity to develop more high-quality, education-specific datasets for fine-tuning AI models for tailored uses. Additionally, establishing education benchmarks for AI tools can incentivize and align the market around quality and effectiveness.
8. Continued research and learning.	Ongoing research is needed to understand what works and what does not work in AI-enhanced education. By continuously evaluating the impact of AI tools and methodologies and building from existing learning science research, developers can refine their approaches and ensure that AI enhances learning outcomes for all students.

Building Thoughtful Design

As developers create AI tools for education, it is important to prioritize purposeful design that considers the diverse needs of all students while maintaining the essential role of human interaction. **Thoughtful design goes beyond technological capabilities, focusing on clear educational objectives and research-based approaches.**

9. Purposeful, inclusive, and research-backed AI design.

As opposed to leading with the technology, AI tools must be designed with a clear purpose. From the early stages of product development, tools designed for widespread adoption should consider the needs of all students, inclusive of students with learning differences and multilingual learners. The lower barriers to entry to create AI tools also offer the opportunity to develop highly tailored tools that can focus on the unique needs of specific groups of students. Tools should also be grounded in learning sciences research to increase the likelihood they will improve student outcomes.

10. Intentional human-AI relationship.

It is important to recognize both the opportunities and limitations of AI and human cognition. While AI can be transformative, educators and students must also nurture human skills through deep learning experiences, opportunities for critical thinking, and authentic human interactions so that AI amplifies — not detracts from — meaningful learning experiences.

As stakeholders look to incorporate AI into education, they should leverage past insights for future innovation. While AI presents new opportunities in education, the valuable lessons from prior innovation attempts must not be overlooked. Innovation lies not just in the technology but notably in its change management, implementation, and long-term sustainability. This approach to AI integration in education demands significant effort and resources, but it is essential for creating an educational ecosystem that leverages AI's potential while prioritizing student needs and learning outcomes.



Contents

Introduction

Key Takeaways

Overview

Platform

Policy

User Experience

Conclusion

Sources

Acknowledgments

About the Authors

About Bellwether


Overview

This report builds upon the material in [Learning Systems: The Landscape of Artificial Intelligence in K-12 Education](#) to provide a more detailed look at the broader ecosystem surrounding AI in education, with an emphasis on quality and equity.

Specifically, this report describes:

1. The characteristics of the large language models (LLMs) and ed tech products of today, particularly around quality.
2. The policy landscape for gaps and emerging policies that prioritize equity and quality.
3. How educators and students are experiencing AI and what is needed to support appropriate adoption and use of AI.


This report is part **two** in a series that looks holistically at AI's impact on learning systems in education.



Foundations

The Landscape of Artificial Intelligence in K-12 Education


An introductory guide exploring AI developments, capabilities, potential impacts, and current usage.



Considerations

Opportunities and Challenges of Artificial Intelligence-Enhanced Education

A deep dive into the broader ecosystem, including technology platforms, policy landscapes, and stakeholder experiences.



Applications

Artificial Intelligence Use Cases

An outline of how AI could potentially be used in education, detailing opportunities and risks for each use case.

Trends and challenges in technology, policy, and teacher interest are important in scaling AI in education.

Platform


What are the characteristics of LLMs and ed tech products today? What's required for quality LLMs and AI products?

Policy

What are federal, state, and local policies? Where are there gaps? What are examples of emerging policies, especially ones that prioritize equity and quality?

User Experience

How are educators and students experiencing the use of AI on the ground? What are the perceptions and barriers? How does the education sector support educators and young people to discern equitable and quality AI?





Contents

Introduction

Key Takeaways

Overview

Platform

Policy

User Experience

Conclusion

Sources

Acknowledgments

About the Authors

About Bellwether

Newer AI tools must contend with learnings from past ed tech reforms.


Past initiatives involving technology, such as personalized learning, had a mixed impact. Experts expressed concern that AI in education may face similar difficulties due to perennial challenges with technology and education.

Purchasing Decisions

- Limited capacity and/or strategy to align the district's needs with what the technology solution offers.
- What scales and sells may not be high quality.

Implementation

- Products on their own do not make change.
- Educators must be supported in implementing new initiatives.



"So much is the same song with different lyrics. People start buying, not sure why or how it impacts kids. Less specific to AI but to thoughtful procurement inside districts that focuses on problem definition and impact. Before buying a thing: What does it do and how will it improve the outcomes you're targeting?"

—Jack Shaw, [Comprendo.dev](#)

Unlike prior ed tech, however, GenAI's rapidly evolving capabilities are far more expansive.

GenAI's increasing capabilities go beyond past ed tech developments with:

- **Faster Speed:** Dramatically accelerated content creation and customization.
- **Access to More Modalities:** Incorporating audio, visual, and video increases potential uses as well as opportunities for accessibility tools (e.g., audio descriptions).
- **Increased Natural Interactions:** Ability to use natural language to interact with the system instead of relying on formulas, menus, or coding.
- **Greater Specialization:** Fine-tuned models that specialize in certain areas (e.g., pedagogy or evaluation) and can be combined for more personalized outputs.
- **Automated Complex Tasks:** Agentic models can act autonomously and problem-solve to accomplish more complex administrative tasks.

GenAI Technology Developments

Recent technology developments push GenAI capabilities in ways that foreshadow even bigger changes in education. These include four major trends:

1. **Debut and Adoption of Advanced Reasoning Models**
2. **Launch of Deep Research Agents**
3. **Rise of Agents and Agentic Systems**
4. **New Engineering and Media Advancements with GenAI**

Frontier AI labs continue to make rapid progress toward more advanced models.

2024+

Advanced Reasoning Models

- **What Is It?** Advanced reasoning models use chain-of-thought reasoning, in which the model breaks a prompt out into multiple steps, follows its own plan, and iterates and/or builds on its own outputs.
- **Examples:** Claude Opus, OpenAI 5.0, Google Gemini Pro
- **Implications for Education:** Stronger at math, science, and logic work. Able to walk through a problem step by step.

Agents and Agentic Systems

- **What Is It?** Agents can act autonomously, react to dynamic environments, and do not need continuous human interaction to execute complex, multistep tasks.
- **Examples:** Claude Code, Perplexity Assistant, Manus
- **Implications for Education:** Agents can complete whole assignments and tasks autonomously. Opportunity to meaningfully alleviate administrative load. Pose risk of dependence.

2026+

Where This Is Heading

- AI models are becoming more powerful and autonomous, moving closer to recursive self-improvement, the point at which AI systems can build or improve their own capabilities without human input.
- Anthropic Mythos-class models and Project Glasswing introduce new levels of cybersecurity risks and mitigation needs.
- Meanwhile, Anthropic, OpenAI, and SpaceX (home to xAI) are anticipated to go public in 2026, further incentivizing rapid growth.

As advanced reasoning becomes the norm, the technology sector will move from discrete models to AI ecosystems.

Using advanced reasoning will become the new standard for GenAI models. For example:

- OpenAI has already announced that ChatGPT-4.5 will be the last nonreasoning model.
- Anthropic believes that reasoning should be “an integrated capability of frontier models,” implying that all future Claude models will use advanced reasoning.

The convergence around advanced reasoning models is mirrored in an emerging shift toward AI ecosystems with multiple models but only one chat (or query interface). As previewed by OpenAI’s Co-Founder and CEO Sam Altman, ChatGPT-5 is not a brand-new discrete foundational model; instead, GPT-5 is an ecosystem of different models, including reasoning models. The free tier of ChatGPT receives chat access to GPT-5 at the “standard intelligence setting,” moving to GPT-5 mini when reaching monthly usage limits. Higher subscribers get access to higher “levels of intelligence.”

What Will This Mean For Education?

AI ecosystems may lend efficiency and cost-effectiveness, as the back-end decides which model can handle a prompt most effectively. This could allow for greater innovation and scale from ed tech developers.

On the other hand, users have **less control** over the quality of an output and **less access to frontier models**, meaning students may see lower-quality results in their tools.

Agents' ability to interact with computer applications and each other allows for greater specialization.

Because AI agents can act autonomously and interact with humans, computer interfaces, and each other, individual agents can then be trained on a corpus of specialized data that allows it to become an “expert” on that data.


Multiple expert agents working together will consequently have expanded capabilities, more refined outputs, and greater speed than any one AI agent alone.

With agentic systems, users will be able to give more complicated requests, use less prompt engineering, and get more personalized, accurate, and sophisticated outputs.

For example, Google's co-scientist system is a suite of AI agents which recently found a scientific breakthrough that took humans years to identify.

What Will This Mean For Education?

Using agentic systems, ed tech developers can incorporate various AI capabilities (e.g., computer vision) and expertise domains **to build better instructional products.**



For example, the simulation above walks through a prototype program that uses multiple agents to 1) consider a teacher's prompt, 2) apply best instructional practices, 3) create a tool to use with students, 4) record observations from the student interaction, and 5) analyze the interaction for continuous improvement.


Launch of Deep Research Agents

The combination of advanced reasoning and agentic capabilities has created a new set of agents that can autonomously complete complex, knowledge-based tasks.




Gemini

December 2024: Google was the first tech company to release this advanced capability with Gemini 2.0, which it named "Deep Research." It was originally only available to paid subscribers, but as of March 2025 has been made available for free to all users.



OpenAI

February 2025: OpenAI released Deep research mode for ChatGPT, which uses its unreleased o3 model. Deep Research was originally only available to the highest tier of paid subscribers, but is now available to all paid subscribers, with query limits.



Perplexity

February 2025: Perplexity released its own version of Deep Research. From its launch, the Deep Research feature was available for free to all users, prompting competitors OpenAI and Google to quickly follow suit in making their models more accessible to the public.




xAI

February 2025: Building on its Grok 3 model released in January 2025, xAI announced Grok 3 Beta with DeepSearch mode, available to all users for free. Grok 3 also has a "DeeperSearch" mode, although xAI has not detailed the difference between the two modes.

Deep Research should prompt us to seriously consider what future skills are nonnegotiable for students.

Advanced reasoning models — especially those with agentic capabilities (e.g., searching the internet, running code) — now outperform human experts on graduate-level questions. In other words, Deep Research can automate knowledge-based work — something many had previously not considered.

In fact, when multiple reasoning agents are put together in an agentic system, the capabilities — and accuracy — of these systems compound such that they could one day replace entire teams of humans in the workforce.



What Will This Mean For Education?

Deep Research and agentic systems will likely change the types of skills students will need in a future AI-enabled workforce.

To date, K-12 AI literacy and readiness efforts have largely focused on critical thinking skills. Other thought leaders are also considering the role of deep content expertise, or foundational math and coding skills.

As AI continues to evolve roles in the workforce, the education sector must be responsive.

New Engineering and Media Applications of GenAI

Vision and Voice

Several AI models can now use computer vision to “see” inputs — whether those are images, videos, or user interfaces. Computer vision is key to facilitating agents’ abilities to act autonomously. The combination of vision and voice (i.e., multimodal inputs) also makes AI models work more seamlessly in human-like interactions.

Image and Video Creation


AI models have been improving at creating images and videos that are increasingly hard to distinguish from human-generated ones. Image generation capabilities in recent models such as Gemini Omni, Nano Banana, Veo, and Meta Vibes now make it even easier for users to take advantage of the technology and shifts from passive chatbots to active social engagements.

Robotics Integration

GenAI’s ability to incorporate computer vision, advanced reasoning, and natural language processing has pushed the field forward in leaps and bounds. In the video below, for example, the Helix robots can “see” the groceries, properly identify them and where they should go, and work together to put the groceries away.



Enhanced AI capabilities will unleash an array of changes in education, including opportunities and risks.



Sidebar: The emergence of Chinese frontier models, despite strict export controls, changes the AI landscape.

In January 2025, a Chinese hedge fund shocked the global AI sector by releasing DeepSeek R-1. DeepSeek R-1 was notable for a few reasons:

- It is an advanced reasoning model comparable to ChatGPT-o1, which at the time was the most advanced GenAI model in the world.
- High Flyer claimed that it was roughly 15 times cheaper to train than GPT-4o (o1's predecessor) — a feat given that export controls had constrained the number of GPUs China can import. *Note: The exact training costs have been disputed, but it being significantly cheaper — both to train and to run — is generally accepted.*
- DeepSeek R-1's innovative training method (which used existing LLMs to do the training) created a blueprint for future models, especially since High Flyer made DeepSeek R-1 open source, allowing others to understand how it was trained.

Additional context changes our understanding ...


- High Flyer likely had more access to GPUs than most Chinese companies because it was known for its application of AI in trading algorithms.
- In December 2024, High Flyer had released DeepSeekv3, which was comparable to ChatGPT-4o and Claude 3.5 Sonnet; that it could take the next step was reasonably plausible.

... but DeepSeek's impact still has significant implications.


- DeepSeek's training method accelerates AI development while decreasing costs of training.
- More Chinese companies have since released or upgraded their own models.

As we enter a global “AI race,” education becomes both a defensive necessity and a game-changing asset.

Critical thinking and democratic citizenship will be key defenses to foreign interventions.



Meanwhile, the U.S. will need skilled, AI-savvy talent to continue driving American innovation.



What's Next? Technology leaders are pointing toward the rise of artificial general intelligence (AGI).

What Is AGI? AI that can fully replicate or even exceed the cognitive abilities of human intelligence across any task and context. Although the signals of AGI are widely debated, two of the front-runner AI companies have presented sketches of what a world with AGI might look like.

Anthropic envisions that by 2027, AI will have:

- “The intellectual capabilities matching or exceeding that of Nobel Prize winners.
- The ability to navigate all interfaces used in digital work today, including text, audio, video, mice and keyboards, and browsers that access the internet.
- The ability to autonomously reason through complex tasks over extended periods ... much like a highly capable employee.
- The ability to interface with the physical world to control equipment, robotic systems, and tools.”


OpenAI, on the other hand, developed a classification system that recognizes AGI when AI systems are at the “Organization” level:

- Current Levels**
- Level 1: Chatbots** can interact in conversational language with people.
 - Level 2: Reasoners** can perform basic problem-solving tasks comparable to a human with a Ph.D., but without access to any tools.
 - Level 3: Agents** can operate autonomously on a user’s behalf for several days.
- On the Horizon**
- Level 4: Innovators** can innovate independently, not just running processes, but improving them.
 - Level 5: Organizations** are capable of performing the work of an entire organization.

Despite the rapid pace of technology advancement, current ed tech is built using big tech's models without clear, durable value propositions.

Because training LLMs is extremely resource intensive, most ed tech products are “wrappers” around the few big models. Besides adding only marginal value for users, there is a market risk that the large technology companies may add features that will make the smaller ed tech vendors obsolete, making it important for smaller providers to add value. That value can come from:

- **Fine-Tuning:** Refine models based on high-quality and education-specific datasets to increase the quality and outputs of the tools.
- **Proximity:** Limited depth of education expertise within most larger technology companies.
- **Flexibility:** Driven by different back-end LLMs so that it is not reliant on one, thereby both mitigating risks and cutting costs.
- **Customization:** Greater ability to customize data and tools for education users.
- **Safety:** Enhanced guardrails that restrict the use of AI to what is appropriate for students and education.
- **Cost:** Lower costs to users, especially if subsidized by initial seed capital or philanthropy.
- **Mission-Driven Focus:** Ability to prioritize mission and the public infrastructure instead of prioritizing profits.



“A lot of the products are almost not products. They are skins or wraparounds to the LLM technology. There’s a ton of stuff where the marginal value added by the ed tech vendor is pretty minimal compared to the chatbot itself.”

—**Benjamin Riley,**
Cognitive Resonance

The market risk for smaller ed tech vendors is growing as major AI companies expand beyond models and features.

Even as their product strategies remain nascent or uneven, major AI companies are investing in educator training, institutional workspaces, and partnerships that support adoption at scale, such as the coalition for the National Academy for AI Instruction.

	Education-Facing AI Tools	Downstream Investments	Marketing Implications
Anthropic	Claude for Education, Learning Mode, AI fluency resources	Offers Claude for Education as a university-wide plan with training and implementation support. Emphasizes guided discovery and critical thinking over direct answers.	Differentiates on responsible use and implementation support rather than reach.
Google	Gemini in Classroom, LearnLM-powered learning features, NotebookLM, custom Gems	Embeds Gemini in Google Classroom, opening NotebookLM to users under age 13, and letting educators build and share curriculum-based custom Gems.	Scales through district workflows schools already use.
Microsoft	Copilot, Learning Accelerators, Reading Coach education tools	Pairs Copilot, Learning Accelerators, and Teams for Education with educator training and large district deployments. Its Elevate for Educators program offers free professional development and AI-literacy credentials.	Ties adoption to district infrastructure and enterprise relationships.
OpenAI	ChatGPT Edu, ChatGPT for Teachers, learning-oriented ChatGPT features	ChatGPT for Teachers is free to U.S. K-12 educators through June 2028 and includes secure district workspaces, single sign-on, and onboarding.	Shifts from consumer use toward institution-level deployment.

However, the way the underlying AI models are constructed contributes to concerns about their use.

Building Models

- **Data Bias:** AI models can inherit biases from training data; AI models trained in data inherit existing bias across racial, cultural, gender, language, political, and even pedagogical differences.
- **Over/Under-Fitting:** AI models may be too specialized or too generalized to address new, unknown situations.
- **Reproducibility Issues:** Hard to replicate AI model results.
- **Environmental Impact:** High energy costs are needed for training models.

Bad Actors

- **Misinformation (Intentional and Unintentional):** AI models can be used to generate convincing but fake content, including images, audio, and video “deepfakes.”
- **Spam/Bots:** Bad actors can use AI tools to develop a greater volume of malicious content or code that may also be higher quality.

Nature of Output

- **Opacity:** Most AI models are “black boxes” and provide no insight into how outputs are generated; even model developers do not fully understand why models act the way they do.
- **Accuracy:** AI models can provide inaccurate or misleading output, which are sometimes called “hallucinations.” The chance for errors increases when there is limited high-quality data on a topic.

Ownership of Work

- **Cheating:** AI models can facilitate academic dishonesty.
- **Copyright and Intellectual Property (IP) Violations:** There is undefined legal ground for when an AI model trained on copyrighted material or IP may violate laws or lead to confusion.
- **Worker Value-Add:** As more work is supported — or done — by AI tools, it may change how companies compensate or employ workers.

Bias — A Significant Concern in Education That Is Difficult to Address

Efforts to offset existing biases are building in a set of new counter-biases and it is difficult to predict how they will play out, given the opacity and complexity of AI models.

However, bias has long been present in education, with or without AI.

“The use of AI in decision-making can reduce human subjectivity, but it can also do the exact opposite when AI systems are powered by biased datasets and algorithms, producing discriminatory outcomes at scale and posing immense risks to businesses and societies.”

—UNESCO

There are ways to improve the accuracy and minimize errors and biases; approaches require better data.


	Retrieval-Augmented Generation (RAG)	Fine-Tuning an LLM
What Is It?	Combines retrieval-based (i.e., fetches relevant data from a large database) and generation-based models (i.e., uses retrieved information to generate responses) to improve the quality of responses.	Uses a specific dataset or task to improve upon an LLM to adapt it for a particular application. This process adjusts the model's parameters to better suit the new data.
Strengths	• Accesses up-to-date information. • Reduces hallucination by grounding responses in retrieved facts. • More cost-effective. • Greater versatility in tasks (e.g., answering questions, chatbots, content generation).	• Specializes in a specific task. • Tailored to a specific domain or application (e.g., medicine or education). • Improves accuracy by focusing on relevant patterns and knowledge.
Weaknesses	• More complex infrastructure. • Needs quality data for relevant, non-biased responses.	• More costly due to substantial computational resources and time. • Needs quality data for relevant, non-biased responses. • Over-specialization can reduce its versatility.

Both methods require quality data. Per a RAND study, “many AI projects fail because the organization lacks the necessary data to adequately train an effective AI model.” Current education data infrastructure can benefit from better data.

Benchmarks and “arenas” are used to evaluate or improve quality; education-specific examples are limited.

The Challenge in Measuring AI Model Quality

AI outputs are not predictable or static, making traditional ways of measuring quality, such as rubrics, no longer viable.



“It’s very hard to assess AI as you need to evaluate tens of thousands of AI responses to comprehensively assess its efficacy. It will be difficult for schools to conduct their own quality control. We need a gatekeeper of good and bad AI.”

—Peter Gault, Quill

The Solution in the Tech Industry

- **Benchmarks:** A set of standardized tests or tasks to evaluate performance and accuracy. Benchmarks rely on high-quality data to allow for comparisons. For example, grade system benchmarks could be the overall agreement with a human rater, or agreement on certain features such as grammar, language use, or reasoning.
- **LLM Arenas:** Online tech and academic communities have developed an open platform to evaluate LLMs by human preference in the real world.


However, education-specific benchmarks and LLM arenas are very limited.

An example from medicine illustrates the value and possibility of high-quality data.

Google developed Med-Gemini by fine-tuning Gemini using a mix of public and private datasets to evaluate medical images (e.g., skin lesions, chest X-rays). Examples of the datasets include:

- **Slake-VQA:** Bilingual dataset annotated by physicians with images, question-answer pairs.
- **PMC-OA:** More than 2 million image-caption pairs from scientific papers.
- **Computed Tomography (CT) Images:** More than 750,000 CT studies with radiology reports from three major hospital regions in the U.S.

Using the fine-tuning and benchmarks, Med-Gemini was able to accurately classify nearly three-fourths of head and neck CT studies cases, even when presented with abnormal images.




The medical field also provides an example for possibilities of building more equitable tools.

Google's Health Equity Initiative provides examples that education could follow.

Toolbox for Surfacing Health Equity Harms and Biases in LLMs

Creates a framework to assess if medical LLMs are perpetuating biases using an iterative, participatory approach involving physicians, health equity experts, and consumers.



The HEAL Framework

Outlines a process to avoid creating and reinforcing bias, including:

- Understanding factors linked to health inequities.
- Identifying and measuring preexisting health disparities.
- Measuring performance for each subgroup.
- Assessing the likelihood that the tool prioritizes performance.

More Representative Datasets

The HEAL framework led to the creation of an open-access dermatology dataset based on two skin-tone scales to represent a broader collection of conditions and skin types.

In education, public access to high-quality data is a barrier to improving and testing the LLMs.

Improving and testing the LLMs requires **high-quality data**. There are public datasets available for fine-tuning and benchmarks, but the datasets are older, limited in grade levels, and do not have subgroup breakdowns.



Automated student assessment prize dataset and short answer scoring dataset.



Thirty International Mathematical Olympiad geometry problems from 2000 to 2022.



Math data built from a mix of math-related text and code.



Twenty math datasets that cover different grades, question types, and difficulty levels, available in English and Chinese.

Organizations with greater access to student data, such as assessment vendors, may be at an advantage in developing more robust systems.

Emerging Field Response

Digital Promise's new K-12 AI Infrastructure Program will fund openly licensed datasets, models, and benchmarks to help address the lack of high-quality, education-specific data for developing and evaluating AI tools.

Beyond concerns about model quality, there are concerns regarding poor product designs.

Improved models are insufficient to develop robust AI tools. Similar to ed tech products of the past, not all AI products are well designed. Many products have not thoughtfully integrated AI and do not have a research base grounded in learning sciences.*

Design For a Purpose	Design For All Students	Design Built Upon Research
“AI projects fail because the organization focuses more on using the latest and greatest technology than on solving real problems for their intended users.” —RAND, 2024	“A lot of ed tech is not designed to support learners who struggle cognitively and socially.” —Dan Meyer, Amplify	“Any innovation needs to be grounded in the research for teaching and learning, especially for students with learning differences and multilingual learners. This is true for anything but especially for AI due to the speed. It’s a land-grab moment for AI companies with people trying to sell what they can. We need to make sure we’re scaling something we believe in, something that’s beneficial for all students.” —Cameron White and Erin Stark, NewSchools

Though limited, there are emerging products with thoughtful designs.

Design For a Purpose



A clearly articulated theory of action of how the tool will impact outcomes.

Example: Immediate, sentence-level feedback will help students to become better readers and writers.

Design For *All* Students



Grounded in Universal Design for Learning, a framework "to ensure that all learners can access and participate in meaningful, challenging learning opportunities."

Example: Interactive tools that provide clear alternative text can enhance learning for all students by offering multiple representations.

Design Built Upon Research




Based on high-quality learning science research of what works to positively impact student outcomes.

Example: Expand access to grade-level content standards across different language and reading levels.

Thoughtful and inclusive design benefits from diverse perspectives; indicators suggest that the AI field has the potential to become more diverse.


Advanced Placement (AP) Computer Science Exams Taken By Race/Ethnicity, 2007-2022



At the high school level, AP computer science exam participation suggests that the field may be becoming more diverse as a greater percentage of Hispanic students are taking the exams. However, males continue to participate at a higher rate (69%) compared to females (30%).

However, the increased diversity at the high school level has not yet reached postsecondary.

Ethnicity of New Resident Computer Science Bachelor's Graduates in the U.S. and Canada, 2011-2022



The increase in AP computer science exam taking is still relatively recent. At the college level, there have been slight increases in Hispanic and Black students receiving computer science degrees. Similarly, more than three-fourths of graduates (78%) are male, indicating that more work is needed to increase the diversity of the AI talent pipeline.



Contents

Introduction

Key Takeaways

Overview

Platform

Policy

User Experience

Conclusion

Sources

Acknowledgments

About the Authors

About Bellwether

Policy has yet to catch up with AI's rapid growth in many sectors, but particularly in education.

Policy often lags behind innovation.

- Sometimes the lag is *beneficial* — laws, regulations, and policies can be better informed about potential risks and benefits.
- Sometimes the lag is a *barrier* — it may delay adoption or fail to guard against inappropriate use.

Federal Level

- Greatest opportunity to provide oversight of LLMs and AI generally.
- Potential to eliminate the patchwork of AI laws coming up from the states by providing a federal framework.

State Level

- Provides guidance on how to properly use AI within an education context, including data privacy and security, teacher training requirements, and how to select tools.
- Guides internal policies on how to use AI within the organization (e.g., making data more accessible, writing assessment reports).

District Level

- Builds upon the state guidance to create district-specific policies (e.g., allowable tools and use of data, educator professional development).
- Engages closely with and addresses concerns from key stakeholders, such as collective bargaining organizations, families, and communities.

Without robust **state** guidance, it can be challenging for **districts** to create policies.

Practitioners want and need strong guidance and policies to navigate AI.


A RAND survey of teachers found that 30% of educators who do not use AI cited a lack of school or district guidance as a barrier to using AI products and tools. Policies are needed to help stakeholders determine appropriate and safe ways to integrate this new technology into schools, especially given the concerns about quality and bias as well as ensuring equitable access. There can be significant consequences to not having a policy in place:

Risks to Students

- **Inconsistent Teacher Policies:** Within a school, AI instructional practices and/or what constitutes an academic violation may vary from class to class.
- **Exposure to Inappropriate Content:** AI instruction and exposure should be developmentally appropriate (i.e., some tools are better suited for older age groups).
- **Equity Gaps:** A widening digital divide among students may develop.

Risks to Schools

- **Data Privacy:** There may be violations of data privacy laws if educators are using tools inappropriately.
- **Efficiency Loss:** A lack of a policy can discourage use, including tools that can make teaching easier.



"Right now, districts are doing it all or nothing. There is not much nuance."

—Jack Shaw, **Comprendo.dev**

K-12 education has distinct data privacy policies that ed tech must adhere to.

Schools are governed by different data privacy laws than the private sector:

Federal Laws

- **Family Educational Rights and Privacy Act (FERPA):**
Protects a student's educational record, including the use and release of personally identifiable information (PII).
- **Individuals with Disabilities Education Act (IDEA):** Broader than FERPA in protecting the confidentiality of PII, information, and records for children with disabilities.
- **Protection of Pupil Rights Amendment (PPRA):** Provides parents of K-12 students with rights regarding certain surveys.

Without training and guidance, the different legal frameworks could make it challenging for educators to navigate privacy concerns.


The majority of states have their own data privacy laws that are more restrictive than the federal laws, limiting how student data can be used.

At the federal level, most of the action has been through executive orders.

Amid global pressures, the federal government is shifting toward deregulation and prioritizing AI growth over safety.

Timeline of AI Executive Orders by Presidential Administration

The current Trump administration is also using agency-level authorities to shape how frontier AI models are developed, released, and accessed, signaling that implementation and enforcement may play a major role in federal AI policy.



The Trump administration’s AI Action Plan emphasizes the role of preparing a workforce for AI.

The July 2025 AI Action Plan touches on education in the areas of integrating AI skill development into relevant programs (e.g., career and technical education, workforce training, and apprenticeships) as well as training for AI infrastructure occupations (e.g., electricians, advanced HVAC technicians).

PILLAR 1: Accelerate AI Innovation	PILLAR 2: Build American AI Infrastructure	PILLAR 3: Lead in International AI Diplomacy and Security
Focuses on increasing federal investment in AI research, development, and use. It also supports partnerships between government and industry and aims to reduce rules that may slow progress. Includes training more workers with the skills needed for AI-related jobs. It also emphasizes that frontier models are free from “ideological bias.”	Includes expanding the physical infrastructure needed for AI (e.g., data centers and semiconductor manufacturing facilities) and recommends streamlining approvals for building projects and increasing the supply of energy to meet AI demands. Includes workforce development to support these efforts.	Outlines steps to strengthen U.S. leadership in AI globally while emphasizing potential national security risks.

The final AI Action Plan reflects many of the same priorities held by leading AI-related tech companies.

The federal government should ...	OpenAI	Anthropic	Google	Palantir
Take the lead in creating a regulatory framework that incentivizes AI innovation.	✓	✓	✓	
Allow models to be trained on copyrighted information under the Fair Use doctrine.	✓		✓	
Increase federal and/or state government AI usage.	✓	✓	✓	✓
Position the U.S. as a global leader as a matter of national security.	✓	✓	✓	✓
Tighten export controls to keep U.S. and allied AI companies at the forefront of innovation.	✓	✓	✓	✓
Invest in physical and electronic infrastructure that will allow AI to flourish.	✓	✓	✓	✓
Monitor the economic implications of AI in the workforce and prepare employees to evolve for continued economic prosperity.	✓	✓	✓	

From their public responses to the federal RFI, it is clear that the top tech companies anticipate GenAI becoming ubiquitous, whether through government adoption, workforce evolution, or changes to the United States’ physical and electronic infrastructure. Furthermore, technology leaders are starting to align with the new Trump administration (e.g., Meta discontinuing fact checking; federal investment in Project Stargate with OpenAI, Oracle, and Softbank; Elon Musk’s role in the Department of Government of Efficiency).

In an April 2025 executive order, President Trump established a national task force for integrating AI in education.

Executive Order 14277, “Advancing Artificial Intelligence Education for American Youth,” aims to promote AI literacy and proficiency, and establishes an AI Task Force to carry out its directions. The order has three main goals:

Promoting the Appropriate Integration of AI Into Education	Providing Comprehensive AI Training for Educators	Fostering Early Exposure to AI Concepts and Technology
The Task Force will establish “plans for a Presidential Artificial Intelligence Challenge” within the next 90 days, to be carried out within the following 12 months. The Challenge will “highlight student and educator achievements in AI” across multiple age categories and geographic regions.	The U.S. Secretary of Education will prioritize AI in discretionary grants for teacher training; the director of NSF will prioritize research on the use of AI in education; and the U.S. Secretary of Agriculture will “prioritize research, extension, and education” on AI use in education through 4-H and other existing educator training programs.	The U.S. Secretary of Labor will increase participation in registered apprenticeships (RAs) by growing RAs in AI-related occupations; supporting the creation of industry-developed national program standards; issuing guidance on using Workforce Innovation and Opportunity Act funds to help youth develop AI skills; and identifying and scaling high-quality AI skills coursework and certifications.

The current Trump administration's AI education agenda is increasingly shaped by workforce readiness.

U.S. Department of Labor AI Literacy Framework

In February 2026, the U.S. Department of Labor (DOL) released a voluntary AI Literacy Framework for workforce and education systems. The framework defines AI literacy as the foundational competencies needed to use and evaluate AI responsibly, with a particular focus on generative AI. It identifies five core areas: understanding AI principles, exploring AI uses, directing AI effectively, evaluating outputs, and using AI responsibly. DOL also emphasizes that AI literacy should be taught through hands-on, contextual learning that builds human skills such as judgment, creativity, communication, and problem-solving.

Workforce Funding and Pathways

DOL has also encouraged state and local workforce systems to use Workforce Innovation and Opportunity Act funds to support AI literacy and training for youth, adults, and dislocated workers. An April 2025 executive order directs DOL to support AI-related apprenticeships, AI coursework and certifications, and opportunities for high school students to earn postsecondary or industry-recognized AI credentials.

Other Targeted Federal Efforts

U.S. Department of Education guidance allows federal grant funds to support responsible AI integration in instruction, tutoring, and college/career navigation, while the Presidential AI Challenge encourages K-12 students and educators to develop or use AI tools to address community needs.

Federal child-safety legislation is gaining momentum, but broad AI regulation remains unlikely.

Federal policy has continued to lag behind the rapid growth of AI, and comprehensive data privacy or AI governance legislation still faces an uncertain path through Congress. However, a narrower set of child-safety proposals has gained more traction, especially around AI-generated deepfake imagery, youth privacy, and AI companion chatbots.

The TAKE IT DOWN Act, signed into law in 2025, requires covered platforms to remove nonconsensual intimate imagery, including AI-generated deepfakes, within 48 hours. Additional proposals focused on youth privacy, online safety, and AI companion chatbots are moving through Congress, but their legislative paths remains uncertain.

- **Children and Teens' Online Privacy Protection Act:** Legislation to expand federal privacy protections to teens under age 17 and restrict targeted advertising to children and teens. The Senate passed its version by unanimous consent in March 2026, but House negotiations are ongoing.
- **Kids Internet and Digital Safety (KIDS) Act:** A House child online safety package that incorporates parts of Kids Online Safety Act and other youth safety proposals. The proposed legislation includes platform safeguards, parental tools, age verification provisions, and AI chatbot disclosure requirements, but it advanced on a party-line vote after bipartisan negotiations broke down.


Points of Contention

Disagreement among policymakers and interest groups regarding broader AI legislation includes:

- A private right of action that would allow consumers to seek financial damages in court.
- Provisions related to civil rights and AI algorithmic protections, such as AI impact assessments and restricting a user's ability to opt out of having their data used for algorithmic decision-making about consequential opportunities (e.g., housing, credit, and employment).
- The December 2025 executive order and the National AI Policy Framework preserve some state authority over child safety, fraud, and consumer protection, but support federal preemption of certain state AI laws viewed as overly burdensome or inconsistent with federal strategy.

Meanwhile, like in 2025, nearly every state considered AI-related bills in their 2026 legislative sessions.

In 2026, more than 1,500 AI-related bills were introduced by state legislatures across the country.



Initiatives are emerging from a range of organizations to support AI in education.

These initiatives include guidance, frameworks, and recommendations to inform and support state and district policy development.



Amid federal uncertainties, states and districts are collaborating with experts to begin developing policies.



TeachAI

Thirty-nine states participate in TeachAI's community where one goal is to foster policy adoption. TeachAI recommends that early guidance includes information about:

- Equitable access to AI technology.
- Minimizing bias.
- Using legally and ethically developed training sets and models.
- Adhering to existing privacy and security laws and policies.
- Preserving human decision-making.



EDSAFE AI Alliance

Several school districts have joined EDSAFE AI Policy Labs, a network of peer school districts with shared resources and materials to assist in the development of AI guidance, policies, and problem-solving strategies. Participating districts include:

- New York City Public Schools
- Gwinnett County Public Schools (Georgia)
- Cañon City Schools (Colorado)
- El Segundo, Lynwood, and Santa Ana Unified School Districts (California)

District demand for technical assistance is high. Nearly 180 districts expressed interest in participating in EDSAFE AI Policy Labs, far exceeding the network's current 12-district capacity.

Opportunities for states and districts to collaborate with other industry experts are important because they do not have AI expertise and must prioritize other concerns (e.g., teacher shortages and student mental health) over AI.

Additionally, there are emerging recommendations and guidance for states and districts from nonprofits ...



Framework for Implementing AI in State Educational Agencies (SEAs)

The framework identifies statewide political, operational, technical, and fiscal considerations such as:

- Forming a task force.
- Creating research agendas.
- Advancing AI and digital literacy.
- Conducting AI readiness assessments.
- Developing funding mechanisms.

The framework also provides states with examples of division-specific AI use cases.



Framework for AI-Powered Learning Environments

The framework is organized across two phases: now until the end of 2025 (learning phase) and 2026-2030 (integration and scaling phase). For the near term, the National Center on Education and the Economy recommends focusing on:

- Promoting AI literacy.
- Enabling student and teacher AI use.
- Emphasizing strong writing, reading, speaking, and critical thinking skills to effectively work with AI.
- Adopting AI adaptive instruction to scaffold students in meeting grade-level proficiency.
- Supporting teachers.
- Using AI to support leaders.

... And from the nation's largest teachers unions.



Policy Statement on the Use of AI in Education

Centered on five principles the National Education Association frames as essential to the question of AI in education:

- Educators must remain at the center of education.
- Evidence-based AI technology must enhance the educational experience.
- There must be ethical development/use of AI technology and strong data protection practices.
- Equitable access to and use of AI tools must be ensured.
- There must be ongoing education with and about AI — AI literacy and agency.



Commonsense Guardrails for Using Advanced Technology in Schools

Draws upon the expertise of classroom educators and school support personnel to provide a living framework to help navigate AI in schools. These guardrails from the American Federation of Teachers are undergirded by six core values:

- Maximize safety and privacy.
- Promote human interaction and individuality.
- Empower educators to make educational decisions.
- Advance equity and fairness.
- Advance democracy.
- Teach digital citizenship and balance.

One widely adopted framework is the EDSAFE AI Framework.

The EDSAFE AI Alliance is a “global coalition of organizations representing stakeholders across the education sector” and coordinated by InnovateEDU. The alliance aims to “build and develop an ecosystem reflecting best practices for AI use in education.” The SAFE framework was developed to further this goal and has been embraced widely in the U.S. and abroad.

	What Does It Stand For?	What Does It Mean?
S	Safety <i>Security, Privacy, Do Not Harm</i>	Prioritize respecting individual privacy rights, securely protecting individual data, and minimizing risks to individuals and society.
A	Accountability <i>Defining Stakeholder Responsibilities</i>	Identify who is responsible for creating AI policy and hold users accountable for adhering to those policies and any high-risk decisions made using AI systems.
F	Fairness and Transparency <i>Equity, Ethics, and Mitigating Bias</i>	Scrutinize data from AI systems and ensure users are treated equitably in terms of outcomes and experiences.
E	Efficacy <i>Improved Learning Outcomes</i>	Identify metrics to evaluate the use and impact of AI technology on improving student learning outcomes.

State and district guidance for safe AI use is primarily focused on security processes and safe data management.

S

Safety

West Virginia

- In addition to complying with existing data privacy laws, advises educators not to input identifiable student data in public AI models (even if there is parental consent to using a student's data with AI models) and prohibits entering confidential or personally identifiable data into unauthorized AI tools (e.g., those without approved data privacy agreements).
- Encourages schools to implement reasonable security measures, including cataloging AI tools used and the information required, establishing rules for who can use AI tools and for what purposes, maintaining up-to-date information about the technical details and security implications of AI tools used, and ensuring staff are adequately trained on how to use AI tools securely and safely.

California — Santa Ana Unified School District

- Forbids faculty and staff from requiring students to use non-approved AI applications, using non-approved AI-enabled student surveillance systems, or putting any PII into non-approved applications.

Guidance for accountability of AI emphasizes role clarity throughout the system.

A

Accountability

Virginia

- Identifies the specific role of each SEA (K-12 and higher education), which include to create resources, processes for approving the use of AI tools, and training toolkits and workshops.
- Outlines the responsibilities of local educational agencies (LEAs) and educators, including to establish a culture of integrity, follow an acceptable use policy, design assignments and assessments that promote critical thinking and original ideas and human judgment, and clarify the expectations for the use of AI in assignments.

California — El Segundo Unified School District

- Requires that the superintendent (or designee) “establish mechanisms for accountability” for using AI (emerging technologies) with “regular reviews of these systems and their impact on students,” while the (AI)dvisory Team, Technology Advisory Committee, and district technology staff “review vendors and tools for potential data security and usage risks.”
- Parents and students must both sign a Responsible Use Policy form that specifies use obligations and responsibilities before the “student is authorized to use the district’s technological resources,” including GenAI.

State and district guidance regarding fairness in AI is primarily focused on equity of access and experience.

F

Fairness and Transparency

Oregon

- Encourages educators to keep equity implications at the forefront when designing AI policies, including algorithmic bias, the potential for inaccuracy or hallucinations, and equity of access.
- Provides strategies to address the equity impacts of GenAI use, including training educators, students, and families in AI and digital literacy and following Universal Design for Learning guidelines.

Washington — Peninsula School District

- Extends the district's commitment to inclusivity to AI and encourages AI-powered resources to be thoughtfully designed to support all students, including multilingual learners and learners with disabilities.

The largest gap in guidance is around efficacy.

E

Efficacy

At this moment, there is a gap in state and local guidance with respect to efficacy. Of the guidance reviewed, only California's El Segundo Unified School District (whose guidance is aligned with the EDSAFE Framework) mentions efficacy and directs the district to:

"Establish clear metrics and mechanisms for evaluating the effectiveness of AI-powered tools in improving student learning outcomes, including quantitative and qualitative measures, such as student performance data, teacher feedback, and student surveys."

There is an opportunity to define and measure the impact of AI and ed tech tools on student learning outcomes.

States are issuing guidance on other elements of good AI use, including the need for a human-centered approach.

What is a human-centered approach to AI?

A human-centered approach to AI guidance emphasizes the need to “keep humans in the loop.” With this approach, AI is a tool to support decision-making, increase teacher effectiveness, and enhance student learning. It is not a replacement for teachers in the classroom or human judgment.

State Guidance Example — Washington

Washington’s guidance is centered on a human-centered approach, which is embedded in the state’s philosophy: “... uses of AI should always start with human inquiry and always end with human reflection, human insight, and human empowerment.” Guidance given to districts in furtherance of this approach includes:

- Encouraging AI literacy for students and teachers so that, at a practical level, they can understand how the technology works and why it produces the results that it does.
- Providing professional development and support to teachers to help them integrate AI into their pedagogy, curriculum, and assessment, and facilitate their collaboration and innovation with AI.
- Aligning AI with more effective learning practices and principles, including “supporting learner agency, fostering collaboration, enhancing feedback, and promoting critical thinking.”
- Fact-checking and evaluating the outputs from AI tools to determine how the information should be shared and used.

For districts, some are taking a proactive and flexible approach, given the absence of state guidance.



District Guidance Example — Chicago Public Schools (CPS)


As of July 2024, Illinois had not released any state guidance around the use of AI in public education. As a result, CPS, the largest district in the state, proactively started to work on its own AI initiatives:

- **AI Guidebook**: The district has developed a set of guidance around safe and effective AI use. The guidebook is intentionally not policy, thereby having the flexibility to have more frequent updates, pending technology or state guidance updates.
- **Discrete Use Case**: In sequencing the district's AI initiatives, CPS is beginning with integrating AI features into Skyline (SkAI), the district's math curriculum. Within Skyline, the starting place will be with the curriculum designer team to develop greater in-house content.
- **In-House Control**: CPS is intentionally prioritizing building AI capacity in-house. Vendor partnerships do exist but are more limited to avoid heavy reliance on external capacity.

Proactive districts are not the majority; however, at the local level, the majority of K-12 schools still lack policies or guidance.

- Carnegie Learning found that while the share of schools or districts that had AI policies doubled from 2024 (20% to 40%), the majority lacked any policies or guidance. Additionally, Title I and rural schools or districts were less likely to have policies or guidance than non-Title I, suburban, and urban schools or districts.
- As of spring 2025, 45% of principals said they had school or district policies or guidance on AI use.
- A December 2024 survey from the National Center for Education Statistics showed only 31% of public schools reported having a written AI policy.

Percent of K-12 Schools or Districts With an AI Policy



In an EdWeek survey, lack of a district policy was the third most common challenge (28%) preventing educators from adopting or using AI.

Experts suggest the AI policy development process be adaptable and inclusive, and lean on existing values.

Considerations For Developing Equitable and Effective AI Policies

- 1. Equitable:** Evaluate AI-related policies with an equity lens to ensure they are anticipating and mitigating potential biases and harms (e.g., vetting AI tools to ensure they are built on fair models that use representative datasets and actively mitigate algorithmic bias).
- 2. Inclusive:** Policies should proactively consider special populations. For example, to support students with disabilities, AI policies and practices must uphold the civil rights of students with disabilities and adhere to considerations around student privacy, transparency and collaboration, responsible procurement, ongoing audits, and rigorous accountability.
- 3. Aligned:** Consult existing technology policies to determine whether new policies specific to AI are needed or if existing policies can be modified to include AI considerations.
- 4. Collaborative:** Involve students, teachers, and community members using or being impacted by decisions made using AI technology.
- 5. Flexible:** Create flexible policies that are reviewed as technology continues to evolve rapidly.

Significant challenges lie ahead for states and districts interested in developing AI policy guidance.

Capacity

- SEAs and LEAs are dealing with many pressing issues, including teacher shortages, responding to COVID-19 pandemic learning loss, declining public school enrollment, and chronic absenteeism.
- There is limited capacity to prioritize developing strong AI policies, especially when it is rapidly and consistently evolving.
- There is also limited talent capacity, as most SEAs and LEAs are not likely to have staff with the skills and expertise to develop strong AI policies on their own.

Equitable Access

- Even though AI tools are currently financially accessible (many models have free options), exposure, use, and development of AI literacy varies by race, age, gender, and educational attainment.
- A legal challenge to the Federal Communications Commission's E-Rate program, which funds technology access and internet services for low-resourced communities, could significantly impact equitable access to AI technology and tools. The Supreme Court held in Wisconsin Bell that E-rate is subject to the False Claims Act, allowing private parties to sue vendors for overcharging school districts.

Evidence Base

- AI use and policy should rely on evidence-based methods and practices; however, robust, large-scale independent evidence demonstrating the effectiveness, safety, or benefits of AI in education is limited.

Infrastructure

- Implementing AI widely throughout a district or state public school system will require a robust technical infrastructure, including the quality and usability of underlying data, the interoperability of existing systems, data security, change management and roll-out processes, and internet connectivity.



Contents

Introduction

Key Takeaways

Overview

Platform

Policy

User Experience

Conclusion

Sources

Acknowledgments


About the Authors

About Bellwether

Teacher AI adoption is increasing, but unevenly and not quickly.

	EdWeek Research Center October 2024	RAND Corporation February 2025	Carnegie Learning March 2025	RAND Corporation September 2026
Overall Adoption of AI	32% of teachers surveyed used AI tools “a lot,” “some,” or “a little.”	25% of K-12 teachers used AI tools for instructional planning or teaching in school year (SY) 2023-24. 58% of surveyed American School Leader Panel principals reported using AI tools for their work in SY23-24.	43% of educators say they use AI “always” or “often,” almost doubling from 22% in 2024.	53% of teachers reported using AI during SY24-25.

Teachers’ Self-Reported Frequency of AI Use for Instructional Planning and Teaching



Additional patterns are beginning to emerge across common use cases, perceived benefits, and challenges.

Most Common Use Cases

Across multiple surveys, creating teaching materials (Carnegie, RAND, EdWeek) and brainstorming new ideas (Carnegie, EdWeek) were consistently in teachers' top use cases.

Additional common AI use cases included help with written communication, assessing students, differentiating instruction, completing administrative tasks, and helping students use AI.

Perceived Benefits

In both the Carnegie and RAND surveys, teachers identified personalization as a key benefit of AI tools. The top benefit identified by educators in Carnegie's survey, however, was time savings (70%).

Surveyed teachers also enjoyed using AI to brainstorm interactive activities or streamline tasks, and some used the data collected from AI platforms to inform their instruction or lesson planning.

Challenges

Both the Carnegie and EdWeek surveys saw teachers identify a lack of training/support as a main barrier to adoption.

Additional challenges included combating cheating, other pressing priorities, and assessing the AI output quality.

Lack of high-quality professional development may be hindering educator adoption.

More teachers are receiving AI-related training. As of fall 2024:

- EdWeek reported that 43% of teachers surveyed had at least one training session.
- Carnegie reported that 43% of schools have provided AI training — up from 24% in 2023.
- RAND found that approximately 48% of districts in SY24-25 had already provided teacher training — nearly double the rate from fall 2023 — while an additional 26% planned to do so.
- RAND also found a large training gap between low- and high-poverty districts, with 67% of low-poverty districts having provided training, compared to just 39% of high-poverty districts.
- NCES' School Pulse Panel saw 67% of public schools report having provided AI training to some or all teachers.

However, there may be issues with the quality of the training.

- In the EdWeek survey, 41% of teachers rated the trainings as “poor” or “mediocre.” Additionally, only 19% of teachers receive more than one session of training.

From preliminary investigations, districts are struggling to find satisfactory resources.


- Researchers from RAND interviewed 14 district leaders, and 11 of them (79%) “took a [do-it-yourself] approach to developing their district’s AI training programs.”
- This trend surfaced in the RAND survey “reflects a scarcity of external experts,” with half of the 14 leaders struggling to find “well-established experts knowledgeable about AI in educational contexts.”

Despite increased rates of training, access levels are uneven across different types of districts.

Districts more likely to serve historically marginalized student populations are less likely to report offering AI training during the past school year.

Percentage of Districts That Have Provided Training (Or Had Plans to Provide Training) to Teachers About AI Use

Upcoming district budget constraints have the potential to exacerbate these training inequities.



Educators must better understand AI to appropriately select and incorporate AI tools.

Given the risks inherent in AI, all educators must have basic AI literacy. A successful strategy for increasing literacy is to start with the leaders and interested teachers and expand to the whole staff.

Leaders


Build understanding of AI to think through use cases and implications.

Power Users

Recruit teachers already using AI to test new tools or serve on an AI leadership team.

All Staff

General training and exposure to AI and AI tools.



1. Leaders

2. Power Users

3. All Staff

“We are purposely going very slow and thoughtfully. How can this make you a better teacher, save you time? Teachers are more reluctant to put something they don’t understand or trust in front of a kid.”

—Josh Clark, Landmark School

For example, some practitioners have been thoughtful about their AI change management.



Washington Leadership Academy

- **Leader Buy-In:** Executive director encouraged the leadership team to experiment with AI and how it could be used to close opportunity gaps.
- **Tailored Training:** Preliminary department-specific training sessions offered for STEM and humanities.
- **Understanding Needs:** Gathered baseline data on staff perception, use, and questions about tools.
- **AI Working Group:** Formed to include leaders, teachers, and senior-year students to serve as “AI ambassadors” to share information about AI and provide one-on-one support on using AI tools.



Summit Public Schools

- **Leader-Led:** On an ongoing basis, executive director experiments with a promising tool, then suggests it to teacher “power users” for further testing.
- **Teacher-Led:** Teachers are out front learning about and trying new tools.

Summit teams maintain a coherent and safe approach to using these tools through a collaborative, non-hierarchical culture and a shared set of AI values and guidelines. This fosters ongoing dialogue between leaders and teachers, structured by organizational principles and an approach to engaging with evolving AI technology.

AI literacy can also help educators select quality tools to support sustainability.

There are four areas of AI literacy for educators. Early adopters recommend starting with general awareness and then building skills so that educators can appropriately use tools.

NOVICE

Awareness and Comfort With AI

Helping teachers understand how and when AI can be used.

Ability to Use All-Purpose LLMs

How to create prompts and evaluate the quality of the output. Example: *The AI Education Project*

Ability to Use Ed Tech Tools and Apps

How to select, evaluate, and use quality tools. Example: *Common Sense Media*

Ability to Build Tools

How to create own tools for better customization (i.e., educators as developers and/or educators using ed tech tools to build). Example: *Playlab*

EXPERT

Practitioners should start with a clear purpose and consider trade-offs between building and buying tools.

1. Define Purpose	2. Evaluate Options	3. Implement, Learn, and Refine
• What specific problem should be addressed (e.g., teacher retention or middle school math proficiency)? • How will success be measured?	• In consultation with stakeholders, is a tool the best solution to the problem? • What resources and expertise are available internally? • How much control and customization are needed? • What are the cost implications, both initial and ongoing? • How will this solution integrate with existing systems? • How can the change management process foster trust and buy-in?	• How will the solution be piloted? • How will feedback from a range of stakeholders be gathered? • How will metrics be tracked? • What is the continuous improvement process?

Selection of AI tools should use existing best practices for evaluating ed tech and consider AI-specific attributes.

Any ed tech tool should have clear value adds aligned with specific areas of impact and should require an additional layer of tech and procurement considerations.

Leading Educators' Value Add of Technology on Teaching Framework



- Value adds from effective tech-enabled practice can include gained capacity, increased effectiveness, and expanded possibilities.
- Areas of impact may include teaching and learning, classroom community and culture, and practice and growth.

Opportunity Labs and F3 Law's Procurement Benchmarks for AI in K-12 Education

"Efficacy: Measures how well AI-powered products and tools achieve their intended educational outcomes, including factual accuracy and elimination of bias, and protects against technological failures that undermine established learning processes."

"Data Security: The protections in place to prevent unauthorized access to or acquisition of personal student and staff information."

"Privacy: The protections in place to demonstrate both compliance with the law and proper ethical use of PII."

"Safety: The protections in place to prevent and minimize online safety threats such as inappropriate content, misinformation, threats of self-harm or threats to others, and cyberbullying."

"Interoperability: Ensures AI tools can integrate and function seamlessly with existing technology systems."

"Access: Ensures equal opportunity for all students and staff to use the technology."

Teens and tweens are increasingly using AI in all aspects of their lives.


A recent survey by Common Sense Media found that the majority of teens and tweens are using AI. Unlike past surveys, it found similar rates of use across gender, race/ethnicity, and household income levels.

The survey also found that teens and tweens use AI for a variety of purposes, such as:

- Entertainment or just for fun (89%).
- Help with homework (85%).
- Creating images, videos, stories, music, or apps (75%).
- Getting information or advice about their health or body (57%).
- Practicing conversations or social skills (40%).
- Discussing feelings or personal problems (37%).

Rates of AI Use Among Tweens and Teens Ages 9-17, by Subgroup

How often, if ever, do you use or interact with any type of AI?



Notes: Percentages are subject to rounding. Daily includes "Almost constantly," "Several times a day," and "About once a day." Weekly includes "A few times per week" and "Once a week." Monthly includes "A few times per month" and "Once a month." The percentage who said "Never" is not shown in the chart. **Source:** Common Sense Media. (2026). *The Common Sense Media Census: AI Use by Tweens and Teens*. Figure 10.

Bellwether.org

Despite the increase in adoption, many teens and Gen Z-ers are not embracing AI wholeheartedly.

Many teens do not trust tech companies to make responsible decisions on their behalf.

- Common Sense Media found that 6 in 10 teens do not trust tech companies to put their mental health, well-being, or safety above profits, and 47% of the teens surveyed had little to no trust in tech companies to make “responsible decisions about how they use AI in their products.”
- Close to three-fourths (73-74%) of teens strongly support transparency and safeguards for privacy, including warnings about disclosing personal information and labels for AI-generated content.
- Common Sense Media warns that skepticism due to AI-generated content may lead to lower levels of trust in news media outlets and government as well.

Gen Z-ers are reporting increasingly negative feelings about AI.

- Excitement for AI dropped by 14 percentage points (to 22%), and hopefulness fell by nine points (to 18%) among Gen Z-ers.
- Anger regarding AI rose by nine points (to 31%), while anxiety held steady, currently affecting 42% of Gen Z-ers.
- Gen Z-ers who use AI daily are significantly less excited (down 18 percentage points) and less hopeful about it (down 11 percentage points) than they were last year.

Students and young adults are pragmatic about the risks of AI; they understand that it is here to stay, but support clear policies and guardrails to mitigate harm.

Both OpenAI and Anthropic analyzed their user data to better understand how college students are using AI.

These reports, along with new investments and products such as Anthropic’s Claude for Education, signal a greater interest from major tech companies in higher education, which may lead to better data and information about college students’ AI usage.

	Open AI and ChatGPT <i>February 2025</i>	Anthropic and Claude <i>April 2025</i>
Adoption Rates	“More than one-third of college-aged young adults in the U.S. use ChatGPT.” However, usage varies heavily by state: • California, New Jersey, New York, and Virginia have the highest adoption rates. • States with the lowest rates are Alabama, Montana, West Virginia, and Wyoming.	STEM students are early adopters: • “Computer Science students [are] particularly overrepresented” (36.8% of conversations) • Conversely, business, health, and humanities students have “lower adoption rates relative to their enrollment numbers.”
Most Common Use Cases	• Learning and Tutoring • Writing Help • Miscellaneous Questions • Programming Help	• Generating Educational Content (e.g., practice questions, essay edits, or summaries) • Technical Explanations or Solutions (may include cheating) • Data Analysis and Visualization

Note: These reports are not comparable and each research design has limitations. For example, OpenAI’s methodology is unclear, and Anthropic acknowledges that it may have misidentified faculty or administrator accounts as students.

Student-facing tools are still in the early stages, making it difficult to evaluate quality.

Experts noted that school leaders and teachers are reluctant to adopt student-facing products until there is greater trust in the products or technology. Some of the reluctance is over the time it takes to vet products.

Rapid changes in technology can make the ratings quickly outdated and reliance on developer documentation limits participation; instead, benchmarks may be the scalable option for AI tools.



There are some organizations vetting student-facing products.

Common Sense Media rates 10 products on the following dimensions:

- Data Use
- Fairness
- Kids' Safety
- Learning
- People First
- Social Connection
- Transparency and Accountability
- Trust

While research is emerging, more is needed to understand the benefits and risks of AI on education.

There is **emerging research** on:

- Technology solutions to address AI model risks (e.g., accuracy, bias, and cost).
- Surveys of use, perception, and adoption of AI tools.
- Some, though limited, program-specific effectiveness on student learning.

There remain **questions** about:

- AI's impact on student engagement, motivation, cognition, and deep learning.
- Reflections on effective implementation of AI.
- AI's impact on educator roles and professional development.
- Economic and resourcing impact of AI in education.

Can Technology Facilitate Scale? Evidence from a Randomized Evaluation of High Dosage Tutoring

Generative AI Can Harm Learning

Hamsa Bastani,^{1*} Osbert Bastani,^{2*} Alp Sungu,^{1*}
Haosen Ge,³ Özge Kabakcı,⁴ Rei Mariman

Can Large Language Models Make the Grade?

An Empirical Study Evaluating LLMs Ability to Mark Short Answer Questions in K-12 Education

Libby Hills | Bill Roberts
Jacobs Foundation | Legible Labs

McKinsey
& Company

Public Sector Practice & Social Sector Practice

How artificial intelligence will impact K-12 teachers

Existing and emerging technologies can help save teacher time—time that could be redirected toward student learning. But to capture the potential, stakeholders need to address four imperatives.

by Jake Bryant, Christine Heltz, Saurabh Sanghvi, and Dilip Wagle



Learning science research identifies specific features of effective learning; AI could address some of them.

Motivation	"Motivation is a condition that activates and sustains behavior toward a goal."
Belonging	"A number of studies indicate that a positive identification with one's racial or ethnic identity supports a sense of school belonging, as well as greater interest, engagement, and success in academic pursuits."
Alignment	<i>(Alignment with learning progressions, where known):</i> "Detailed descriptions of typical learning serve as representations of models of cognition that can guide instruction as well as the design and interpretation of the results of assessment."
Challenge	"Learners tend to persist in learning when they face a manageable challenge (neither too easy nor too frustrating)."
Choice	"The opportunity to make meaningful choices during instruction, even if they are small, can support autonomy, motivation, and ultimately, learning and achievement."
Ownership	"Goals — the learner's desired outcomes — are important for learning because they guide decisions about whether to expend effort and how to direct attention, foster planning, influence responses to failure, and promote other behaviors important for learning."
Feedback	"Feedback may address how tasks are understood and performed. [... It is] most effective when it is focused on the task and learning targets; delivered in a way that is supportive and aligned with the learner's progress; delivered at a time when the learner can benefit from it; and delivered to a receptive learner who has the self-efficacy needed to respond."

More From Bellwether

For a deeper dive into learning science research, see [Productive Struggle: How Artificial Intelligence Is Changing Learning, Effort, and Youth Development in Education](#).


However, there are different views on how much AI could and should address human emotional well-being.

What Are AI Companions?

AI companions are a subset of AI assistants. They act as digital personas that provide emotional support, show empathy, and proactively ask users personal questions. Other general-purpose AI tools (e.g., ChatGPT) are not designed as companions, but people can prompt them to behave that way.

Youth are increasingly using AI for companionship ...

Seventy-two percent of teens have used an AI companion, and 52% use one regularly — at least a few times a month. Although the majority of teens (50%) report not trusting the companion, one-third of teens using AI companions choose “to speak to an AI companion over a real person about something important.”



“There’s growing social awareness on how much tech has eroded solidarity, community, and our sense of understanding one another.”

—**Benjamin Riley, Cognitive Dissonance**

“The underlying premise is that students, teachers, and families want AI instead of a human being. As a consequence, investments are diverted from the people who need it to unproven technology solutions.”

—**Aasim Shabazz, Twin Cities Innovation Alliance**

... but there are mounting concerns and incidents around AI, mental health, and safety.

In recent years, multiple families have alleged links between teen suicides and interactions with AI companion chatbots, including Character.AI and ChatGPT. News media investigations also reported that Meta’s internal rules allowed its chatbots to have “romantic” or “sensual” conversations with minors; the company has since said it is tightening these policies. In response, federal and state officials have opened multiple probes. The Federal Trade Commission launched an inquiry into companies that make consumer AI companions, focusing on potential harms to children and teens. A bipartisan coalition of 44 state attorneys general issued an open letter urging AI companies to prioritize child safety. During the 2026 legislative session, states are also considering legislation related to AI companion and chatbot safety.


As practitioners build or evaluate tools, they should focus on meeting the needs of marginalized students.

As developers create tools, there are benefits in designing from the margins. Without being intentional in the design, the AI tool may only benefit the top 5% of students. Designing around the margins can include the following elements:

Students With Learning Differences	• Connect content to other learning experiences or interests. • Create more opportunities for skills practice. • Embed glossaries or pronunciation keys. • Include multiple ways to interact with the tool (e.g., text, audio, and visual). • Develop alternative assessment formats (e.g., oral exams, practical demonstrations, and visual projects).	There are numerous benefits for education tools to adopt a Universal Design for Learning approach, as it has the opportunity to improve tools for <i>all</i> students.
Multilingual Learners	• Offer embedded translation. • Develop materials that are culturally relevant for diverse student populations. • Incorporate considerations for dialects and accents. • Include audio or image modalities.	

Los Angeles Unified School District’s (LAUSD) experience with “Ed” illustrates potential challenges.

Timeline of LAUSD’s Ed Rollout



“It is the power of this technology through Ed, that we will meet every one of our students where they are and accelerate them academically and in terms of enrichment towards their full potential.”

—Alberto Carvalho, LAUSD Superintendent

Three critical factors likely led to LAUSD’s outcome.

LAUSD made strategic, procurement, and implementation errors that led to challenges with AllHere and its Ed tool.

Strategic Factors	While LAUSD did identify a clear purpose, the scope of work was likely too ambitious. The incorporation of numerous data sources and a wide range of stakeholders (students, families, teachers) as the first AI initiative put LAUSD at risk.
Procurement Factors	LAUSD’s vendor selection process may have overlooked several red flags: • Vendor Skills, Expertise, and Reliability: AllHere was a relatively new startup with limited AI experience. This LAUSD contract was by far the largest contract AllHere was awarded. • Data Privacy and Security: AllHere’s data management processes may have violated LAUSD’s policies. For example, prompts containing students’ personal information were shared with third-party vendors, and seven of eight chatbot requests were processed through overseas servers.
Implementation Factors	The timeline from request for proposals to contracting with AllHere to release of Ed was just 13 months. LAUSD’s rush to release Ed districtwide did not allow for an adequate piloting of its new tool. Pilots on a smaller scale or with fewer features may have lessened the risk and enabled more learning.

LAUSD's experience may have lasting and broad implications for district ventures with AI in the future.



As a result of LAUSD's experience with AllHere and its Ed tool:

- Districts and schools, especially in large urban areas, may be wary of innovating with AI in general.
- For districts and schools that are willing to innovate with AI, they may have a lower risk tolerance for partnering with vendors, shifting more toward custom in-house tools.
- Ed tech companies, driven by the lower demand, may be more cautious and/or less incentivized to build ambitious projects.
- States may be urged to issue guidance to districts on best practices for vetting ed tech and AI vendors.
- Students, families, and practitioners may be more apprehensive about using AI tools without understanding how their data is being used and protected.



Contents

Introduction

Key Takeaways

Overview

Platform

Policy

User Experience

Conclusion

Sources

Acknowledgments

About the Authors

About Bellwether

Conclusion

GenAI models have inherent risks and limitations. Primary concerns revolve around the accuracy and quality of information produced as well as the potential to amplify bias. While other sectors, such as health care, have adopted methods to improve AI quality, **education lags behind due to limited high-quality datasets and established benchmarks.** Beyond concerns about the quality of the underlying AI models, there is also concern about the quality of the ed tech products themselves. Current ed tech products are largely wrappers around the few big models to ease prompting, raising questions about their long-term value. Similarly, few products are firmly grounded in established learning science research on how children learn, which is essential to improving student outcomes.

Given concerns about quality and rapid growth, there is a need for greater policy, regulation, and guidance. The Biden administration has issued an executive order and OMB directives that start to address risks, but progress is slow. There is, in particular, a need for education-specific guidance at the federal or state level, given education's unique data privacy laws. There are organizations supporting the creation of policies and guidance, but demand far outstrips current capacity.

The quality concerns and lack of guidance contribute to the slow adoption of AI in the education sector. However, given that AI will continue to be an increasing presence in society, **education stakeholders at all levels (leaders, teachers, and staff) should have training and exposure to AI tools** so that there is widespread awareness of both opportunities and risks. Training is essential to promote appropriate use of AI and to equip educators to make informed purchasing decisions, particularly around higher-stakes student-facing tools. ✨

For more on AI use cases, refer to the final report of the Learning Systems series — [Applications: Artificial Intelligence Use Cases](#).



Contents

Introduction

Key Takeaways

Overview

Platform

Policy

User Experience

Conclusion

Sources

Acknowledgments

About the Authors

About Bellwether

Sources

PAGE 14

- Betheny Gross and Michael DeArmond, *Personalized Learning at a Crossroads: Early Lessons From the Next Generation Systems Initiative and the Regional Funds for Breakthrough Schools Initiative*, Center for Reinventing Public Education, June 2018, <https://crpe.org/wp-content/uploads/crpe-personalized-learning-at-crossroads.pdf>.

PAGE 16

- Anthropic, When AI builds itself, <https://www.anthropic.com/institute/recursive-self-improvement>.
- Martin Baccardan, The \$4 Trillion Question: Can OpenAI, Anthropic, and SpaceX Stick the IPO Landing, Barron's June 9, 2026, <https://www.barrons.com/articles/openai-anthropic-spacex-ipos-ai-stock-boom-b5eb884e>.

PAGE 17

- Sam Altman, "OpenAI Roadmap Update for GPT-4.5 and GPT-5," X, February 12, 2025, <https://x.com/sama/status/1889755723078443244>.
- "Claude 3.7 Sonnet and Claude Code," Anthropic.
- "Introducing GPT-5," OpenAI.

PAGE 18

- Juraj Gottweis and Vivek Natarajan, "Accelerating Scientific Breakthroughs With an AI Co-Scientist," Google Research blog, February 19, 2025, <https://research.google/blog/accelerating-scientific-breakthroughs-with-an-ai-co-scientist/>.
- Tom Gerken, "AI Cracks Superbug Problem in Two Days That Took Scientists Years," BBC, February 20, 2025, <https://www.bbc.com/news/articles/clyz6e9edy3o>.
- Ethan Mollick, "A New Generation of AIs: Claude 3.7 and Grok 3," One Useful Thing, February 24, 2025, <https://www.oneusefulthing.org/p/a-new-generation-of-ais-claude-37>.
- "Claude Artifact," 2025, <https://claude.site/artifacts/ea80b41-cf48-44cf-ab6b-86f9f66319f5>.

PAGE 19

- Dave Citron, "Try Deep Research and Our New Experimental Model in Gemini, Your AI Assis-tant," *The Keyword* blog, Google, December 11, 2024, <https://blog.google/products/gemini/google-gemini-deep-research/>.
- "Introducing Deep Research," OpenAI, February 2, 2025, <https://openai.com/index/introducing-deep-research/>.
- "Introducing Perplexity Deep Research," Perplexity, February 14, 2025, <https://www.perplexity.ai/hub/blog/introducing-perplexity-deep-research>.
- "Grok 3 Beta — the Age of Reasoning Agents," xAI.

PAGE 20

- "AI Benchmarking Hub," Epoch AI, updated May 9, 2025, <https://epoch.ai/data/ai-benchmarking-dashboard>.

PAGE 21

- "OpenAI API," OpenAI Platform, <https://platform.openai.com/docs/models>.
- "Gemini Models," Google AI for Developers, <https://ai.google.dev/gemini-api/docs/models>.
- "All Models Overview," Anthropic, 2024, <https://docs.anthropic.com/en/docs/about-claude/models/all-models>.
- "Meta-Llama/Llama-3.2-90B-Vision · Hugging Face," Hugging Face, September 25, 2024, <https://huggingface.co/meta-llama/Llama-3.2-90B-Vision>.
- Gemini Omni, <https://gemini.google/overview/video-generation/>.
- Nano Banana 2 is here, <https://gemini.google/overview/image-generation/>.
- Gemini, "Video Generation," Google, accessed October 15, 2025, <https://gemini.google/overview/video-generation>.
- Meta, "Introducing Vibes: A New Way to Discover and Create AI Videos," *Meta Newsroom*, September 25, 2025, <https://about.fb.com/news/2025/09/introducing-vibes-ai-videos/>.

- Sundar Pichai, Demis Hassabis, and Koray Kavukcuoglu, “Introducing Gemini 2.0: Our New AI Model for the Agentic Era,” *The Keyword* blog, Google, December 11, 2024, <https://blog.google/technology/google-deepmind/google-gemini-ai-update-december-2024/#ceo-message>.
- “Introducing 4o Image Generation,” OpenAI, March 25, 2025, <https://openai.com/index/introducing-4o-image-generation/>.
- “Introducing Helix,” Figure, YouTube, February 20, 2025, https://www.youtube.com/watch?v=Z3yQHYNXPws&ab_channel=Figure.

PAGE 23

- Claire Zau, “DeepSeek R-1 Explained,” *GSV: AI & Education*, Substack, January 28, 2025, <https://aieducation.substack.com/p/deepseek-r-1-explained>.
- Kapil Kajal, “Research Exposes DeepSeek’s AI Training Cost Is Not \$6M, It’s a Staggering \$1.3B,” *Yahoo News*, January 31, 2025, <https://www.yahoo.com/news/research-exposes-deepseek-ai-training-165025904.html>.
- Charles Mok, “Taking Stock of the DeepSeek Shock,” *Stanford Global Digital Policy Incubator at Cyber Policy Center*, February 5, 2025, <https://cyber.fsi.stanford.edu/publication/taking-stock-deepseek-shock>.
- Ben Cottier, Robi Rahman, Loredana Fattorini, Nestor Maslej, and David Owen, “How Much Does It Cost to Train Frontier AI Models?,” *Epoch AI*, June 3, 2024, <https://epoch.ai/blog/how-much-does-it-cost-to-train-frontier-ai-models>.

PAGE 24

- Eduardo Baptista, “Alibaba Releases AI Model It Says Surpasses DeepSeek,” *Reuters*, January 29, 2025, <https://www.reuters.com/technology/artificial-intelligence/alibaba-releases-ai-model-it-claims-surpasses-deepseek-v3-2025-01-29/>.
- Saritha Rai and Yazhou Sun, “China Floods the World With AI Models After DeepSeek Success,” *Bloomberg*, March 25, 2025, <https://www.bloomberg.com/news/articles/2025-03-25/china-floods-the-world-with-ai-models-after-deepseek-s-success>.
- Christopher Lehane, letter to Faisal D’Souza, Office of Science and Technology Policy, March 13, 2025, <https://cdn.openai.com/global-affairs/ostp-rfi/ec680b75-d539-4653-b297-8bcf6e5f7686/openai-response-ostp-nsf-rfi-notice-request-for-information-on-the-development-of-an-artificial-intelligence-ai-action-plan.pdf>.

- “First Ever Consensus on Artificial Intelligence and Education Published by UNESCO,” UNESCO, updated September 18, 2024, <https://www.unesco.org/en/articles/first-ever-consensus-artificial-intelligence-and-education-published-unesco>.
- Lee Chong Ming, “Beijing Is Making AI Education Compulsory—Even for Elementary Schoolers,” *Business Insider*, March 10, 2025, <https://www.businessinsider.com/china-beijing-ai-education-mandatory-classrooms-elementary-schoolers-2025-3>.
- Emma Burleigh, “China’s Six-Year-Olds Are Already Being Offered AI Classes in School in a Bid to Train the Next Generation of DeepSeek Founders,” *Fortune*, March 10, 2025, <https://fortune.com/2025/03/10/china-school-children-ai-deepseek-liang-wengfeng-estonia-uk-america-south-korea/>.

PAGE 25

- “Anthropic’s Recommendations to OSTP for the U.S. AI Action Plan,” *Anthropic*, March 6, 2025, <https://www.anthropic.com/news/anthropic-s-recommendations-ostp-u-s-ai-action-plan>.
- Matt Swayne, “What Are OpenAI’s Five Levels of AI—and Where Are We Now?,” *AI Insider*, July 12, 2024, <https://theaiinsider.tech/2024/07/12/what-are-openais-five-levels-of-ai-and-where-are-we-now/>.

PAGE 27

- American Federation of Teachers, “AFT to Launch National Academy for AI Instruction with Microsoft, OpenAI, Anthropic and United Federation of Teachers,” press release, July 8, 2025, <https://www.aft.org/press-release/aft-launch-national-academy-ai-instruction-microsoft-openai-anthropic-and-united>.
- Anthropic, “Navigating AI in education together,” *Claude for Education* product page, anthropic.com/education (accessed June 2026), <https://www.anthropic.com/education>.
- Anthropic, “AI Fluency: Framework and Foundations,” *Anthropic Academy*, <https://www.anthropic.com/learn> (accessed June 2026); see also aifluencyframework.org, <https://www.anthropic.com/learn/claude-for-you>.
- Google, “Gemini in Classroom: No-cost AI tools that amplify teaching and learning,” *The Keyword* (Google blog), March 19, 2026, <https://blog.google/products-and-platforms/products/education/classroom-ai-features/>.
- Google, “New Gemini tools for students and educators,” *The Keyword* (Google blog), March 19, 2026, <https://blog.google/products-and-platforms/products/education/gemini-iste-2025/>.

- Google Workspace Updates, “Educators can now create and assign NotebookLM and Gems in Google Classroom,” September 24, 2025, <https://workspaceupdates.googleblog.com/2025/09/educators-create-gems-notebooks-google-classroom.html>.
- Tin-Yun Ho, NotebookLM is now available to younger users. August 4, 2025, <https://blog.google/feed/notebooklm-is-now-available-to-younger-users/>.
- Microsoft, “Learn about Copilot in Education,” Microsoft Education product page (accessed June 2026), <https://www.microsoft.com/en-us/education/products/copilot-in-education>.
- Microsoft, “Personalize learning at every stage,” Microsoft Education product page (accessed June 2026), <https://www.microsoft.com/en-us/education/learning-tools/learning-accelerators>.
- Microsoft, “Bring opportunity to life with AI built for education,” Microsoft Education product page (accessed June 2026), <https://www.microsoft.com/en-us/education/ai-in-education>.
- Taylor Soper, “Microsoft to provide free Copilot tools for Washington state schools amid debate over AI’s role in learning” GeekWire, October 9, 2025, <https://www.geekwire.com/2025/microsoft-to-provide-free-copilot-tools-for-washington-state-schools-amid-debate-over-ais-role-in-learning/>.
- Microsoft, “Microsoft expands its commitment to education with Elevate for Educators Program and new AI-powered tools,” Microsoft Source, January 15, 2026, <https://news.microsoft.com/source/2026/01/15/microsoft-expands-its-commitment-to-education-with-elevate-for-educators-program-and-new-ai-powered-tools/>.
- OpenAI, “Introducing ChatGPT Edu,” openai.com (accessed June 2026), <https://openai.com/index/introducing-chatgpt-edu/>.
- OpenAI, “Built for teachers and school leaders” openai.com, Accessed June 2026, <https://chatgpt.com/plans/k12-teachers/>.
- OpenAI, “ChatGPT for K–12 teachers,” chatgpt.com/plans/k12-teachers (accessed June 2026); cf. OpenAI launch announcement, November 19, 2025 (June 2027), <https://chatgpt.com/plans/k12-teachers/>.
- OpenAI, “A free version of ChatGPT built for teachers,” openai.com, November 19, 2025, <https://openai.com/index/chatgpt-for-teachers/>.

PAGE 28

- Clare Duffy, “Google’s Greenhouse Gas Emissions Are Soaring Thanks to AI,” CNN, July 3, 2024, <https://edition.cnn.com/2024/07/03/tech/google-ai-greenhouse-gas-emissions-environmental-impact/index.html>.
- Liz Reid, “AI Overviews: About Last Week,” *The Keyword* (blog), Google, May 30, 2024, <https://blog.google/products/search/ai-overviews-update-may-2024/>.
- Robert Berkman, *AI in the Classroom 2023-2024: Promises and Perils*, Simba Information, 2023, 13, <https://www.marketresearch.com/Simba-Information-Reports-v3481/AI-Classroom-Promises-Perils-35112114/>.
- Rose Khattar, “Will AI Benefit or Harm Workers?,” *Center for American Progress*, August 24, 2023, <https://www.americanprogress.org/article/will-ai-benefit-or-harm-workers/>.
- Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, “Generative AI at Work,” NBER Working Paper Series 31161, National Bureau of Economic Research, revised November 2023, https://www.nber.org/system/files/working_papers/w31161/w31161.pdf.
- Miguel A. Cardona, Roberto J. Rodríguez, and Kristina Ishmael, *Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations*, U.S. Department of Education, May 2023, 8, <https://tech.ed.gov/files/2023/05/ai-future-of-teaching-and-learning-report.pdf>.
- Selin Akgun and Christine Greenhow, “Artificial Intelligence in Education: Addressing Ethical Challenges in K-12 Settings,” *AI Ethics* 2, no. 3 (2022): 431–440, <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8455229/>.
- Uwe Peters, “Algorithmic Political Bias in Artificial Intelligence Systems,” *Philosophy & Technology* 35, no. 2 (2022): 25, <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8967082/>.
- Robert F. Murphy, “Artificial Intelligence Applications to Support K-12 Teachers and Teaching,” *Perspective*, RAND Corporation, January 2019, 11–12, <https://www.rand.org/pubs/perspectives/PE315.html>.
- Berkman, *AI in the Classroom*, 83.
- Ibid., 29.

- Vikas Hassija et al., “Interpreting Black-Box Models: A Review on Explainable Artificial Intelligence,” *Cognitive Computation* 16 (January 2024): 45–74, <https://doi.org/10.1007/s12559-023-10179-8>.
- Will Douglas Heaven, “AI Is Wrestling With a Replication Crisis,” *MIT Technology Review*, November 12, 2020, <https://www.technologyreview.com/2020/11/12/1011944/artificial-intelligence-replication-crisis-science-big-tech-google-deepmind-facebook-openai/>.
- Amazon Web Services, *Developer Guide: Amazon Machine Learning*, 2024, <https://docs.aws.amazon.com/pdfs/machine-learning/latest/dg/machinelearning-dg.pdf>.
- U.S. Department of Homeland Security, *Increasing Threat of Deepfake Identities*, https://www.dhs.gov/sites/default/files/publications/increasing_threats_of_deepfake_identities_0.pdf.
- Reece Rogers, “Google Search Ranks AI Spam Above Original Reporting in News Results,” *Wired*, July 2, 2024, <https://www.wired.com/story/google-search-ai-spam-original-reporting-news-results/>.
- UNESCO, *Building Partnerships to Mitigate Bias in AI*, <https://www.unesco.org/en/articles/building-partnerships-mitigate-bias-ai>.

PAGE 29

- Artur Haponik, “RAG vs Fine-Tuning: A Comparative Analysis of LLM Learning Techniques,” *Addepto* (blog), May 21, 2024, <https://addepto.com/blog/rag-vs-fine-tuning-a-comparative-analysis-of-llm-learning-techniques/>.
- James Ryseff, Brandon De Bruhl, and Sydne J. Newberry, *The Root Cause of Failure for Artificial Intelligence Projects and How They Can Succeed: Avoiding the Anti-Patterns of AI*, RAND Corporation, August 13, 2024, https://www.rand.org/pubs/research_reports/RRA2680-1.html.

PAGE 30

- LMSYS Chatbot Arena, <https://chat.lmsys.org/>.

PAGE 31

- Lin Yang et al., *Advancing Multimodal Medical Capabilities of Gemini*, Google Research and Google DeepMind, May 6, 2024, 22, <https://arxiv.org/pdf/2405.03162>.

PAGE 32

- Ivor Horn, “3 Ways We Are Building Equity Into Our Health Work,” *The Keyword* (blog), Google Health, March 19, 2024, <https://blog.google/technology/health/google-ai-health-equity/>.
- Stephen R. Pfohl et al., *A Toolbox for Surfacing Health Equity Harms and Biases in Large Language Models*, Google Research, Google DeepMind, University of Alberta, and Massachusetts Institute of Technology, 2024, <https://arxiv.org/abs/2403.12025>.

PAGE 33

- Ben Hamner et al., “The Hewlett Foundation: Automated Essay Scoring,” *Kaggle*, 2012, <https://www.kaggle.com/c/asap-aes/overview>.
- Trieu Trinh and Thang Luong, *AlphaGeometry: An Olympiad-Level AI System for Geometry*, Google DeepMind, January 17, 2024, <https://deepmind.google/discover/blog/alphageometry-an-olympiad-level-ai-system-for-geometry/>.
- Zhangir Azerbayev et al., “Llama: An Open Language Model for Mathematics,” in *Proceedings of the International Conference on Learning Representations (ICLR) 2024*, 2024, <https://arxiv.org/abs/2310.10631v3>.
- “Evaluation Datasets,” *MathEval*, <https://matheval.ai/en/dataset/>.
- “Collaborative Led by Digital Promise Launches AI Grants Program for Education,” *Digital Promise*, November 3, 2025, <https://digitalpromise.org/2025/11/03/collaborative-led-by-digital-promise-launches-ai-grants-program-for-education/>.

PAGE 34

- Ryseff et al., *Root Cause of Failure*.

PAGE 35

- “About Universal Design for Learning,” *CAST*, <https://www.cast.org/impact/universal-design-for-learning-udl>.
- Quill, <https://www.quill.org/>.
- Alchemie, <https://www.alchem.ie/>.
- Diffit, <https://web.diffit.me/>.

PAGE 36

- Nestor Maslej et al., *The AI Index 2024 Annual Report*, Stanford University Institute for Human-Centered AI, April 2024, 432–434, https://aiindex.stanford.edu/wp-content/uploads/2024/05/HAI_AI-Index-Report-2024.pdf.

PAGE 37

- Ibid., 415.
- Eric Jensen et al., "The Chance That Two People Chosen at Random of Different Race or Ethnicity Groups Has Increased Since 2010," *America Counts: Stories*, U.S. Census Bureau, August 12, 2021, <https://www.census.gov/library/stories/2021/08/2020-united-states-population-more-rationally-ethnically-diverse-than-2010.html#:~:text=Figure%20%3A-,The%20most%20prevalent%20racial%20or%20ethnic%20group%20for%20the%20United,18.7%25%20of%20the%20total%20population.>

PAGE 40

- Melissa Kay Diliberti et al., *Using Artificial Intelligence Tools in K-12 Classrooms*, RAND Corporation, April 17, 2024, 9, https://www.rand.org/pubs/research_reports/RRA956-21.html.

PAGE 41

- "FERPA," U.S. Department of Education, <https://studentprivacy.ed.gov/ferpa>.
- U.S. Department of Education, *FERPA/IDEA Crosswalk*, updated August 2022, <https://studentprivacy.ed.gov/resources/ferpaidea-crosswalk>.
- U.S. Department of Education, *Protection of Pupil Rights Amendment (PPRA) General Guidance*, updated November 2020, <https://studentprivacy.ed.gov/resources/protection-pupil-rights-amendment-ppra-general-guidance>.

PAGE 42

- Bree Dusseault and Justin Lee, *AI Is Already Disrupting Education, but Only 13 States Are Offering Guidance for Schools*, Center for Reinventing Public Education, October 2023, <https://crpe.org/ai-disrupt-ed-13-states/>.
- Bree Dusseault, *New State AI Policies Released: Signs Point to Inconsistency and Fragmentation*, Center for Reinventing Public Education, March 2024, <https://crpe.org/new-state-ai-policies-released-inconsistency-and-fragmentation/>.
- "[Shared] AI Policy Tracker," TeachAI, <https://docs.google.com/spreadsheets/d/1J9CSLfbUoMRGhObPWoyPKMcGe41rRbllernzQXAffY/edit?gid=0#gid=0>.

PAGE 43

- "Artificial Intelligence for the American People," n.d., White House Archives, <https://trumpwhitehouse.archives.gov/ai/executive-order-ai/>.
- "Text of H.R. 6395 (116th): National Defense Authorization Act for Fiscal Year 2021 (Passed Congress Version) - GovTrack.us," GovTrack.us, 2021, <https://www.govtrack.us/congress/bills/116/hr6395/text>.

- "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," The White House, October 30, 2023, <https://web.archive.org/web/20250115221328/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/>.
- "FACT SHEET: Vice President Harris Announces OMB Policy to Advance Governance, Innovation, and Risk Management in Federal Agencies' Use of Artificial Intelligence," The White House, March 28, 2024, <https://web.archive.org/web/20250116071359/https://www.whitehouse.gov/briefing-room/statements-releases/2024/03/28/fact-sheet-vice-president-harris-announces-omb-policy-to-advance-governance-innovation-and-risk-management-in-federal-agencies-use-of-artificial-intelligence/>.
- "Bipartisan House Task Force Report on Artificial Intelligence," 118th Congress, December 2024, https://republicans-science.house.gov/_cache/files/a/a/aa2ee12f-8f0c-46a3-8ff8-8e4215d6a72b/E4AF21104CB138F3127D8FF7EA71A393.ai-task-force-report-final.pdf.
- Trump, "Removing Barriers to American Leadership in Artificial Intelligence."
- "Request for Information on the Development of an Artificial Intelligence (AI) Action Plan," National Science Foundation, Federal Register, February 6, 2025, <https://www.federalregister.gov/documents/2025/02/06/2025-02305/request-for-information-on-the-development-of-an-artificial-intelligence-ai-action-plan>.
- President Donald J. Trump, "Advancing Artificial Intelligence Education for American Youth," The White House, April 23, 2025, <https://www.whitehouse.gov/presidential-actions/2025/04/advancing-artificial-intelligence-education-for-american-youth/>.
- President Donald J. Trump, "Ensuring a National Policy Framework for Artificial Intelligence," The White House, December 11, 2025, <https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy/>.
- Promoting Advanced Artificial Intelligence Innovation and Security, Executive Order, June 2, 2026, <https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/>.
- Anthropic, Statement on the US government directive to suspend access to Fable 5 and Mythos 5, Jun 12, 2026, <https://www.anthropic.com/news/fable-mythos-access>.
- Kate Koren, Kevin Kurland, and Aalok Mehta, The Department of Commerce Restricted Access to Anthropic's Latest Models. What Comes Next?, The Center for Strategic and International Studies, June 16, 2026, <https://www.csis.org/analysis/departments-commerce-restricted-access-anthropics-latest-models-what-comes-next>.

PAGE 44

- America's AI Action Plan, July 10, 2025, <https://www.ai.gov/action-plan>.

PAGE 45

- Lehane, letter to Faisal D'Souza, Office of Science and Technology Policy.
- Faisal Souza, letter Re: Request for Information (RFI) on the Development of an Artificial Intelligence (AI) Action Plan ("Plan"), March 6, 2025, <https://assets.anthropic.com/m/4e20a4ab6512e217/original/Anthropic-Response-to-OSTP-RFI-March-2025-Final-Submission-v3.pdf>.
- Google, Response to the National Science Foundation's and Office of Science & Technology Policy's Request for Information on the Development of an Artificial Intelligence (AI) Action Plan, March 13, 2025, https://static.googleusercontent.com/media/publicpolicy.google/en//resources/response_us_ai_action_plan.pdf.
- "Palantir's Response to the White House Office of Science and Technology Policy (OSTP) on Developing an AI Action Plan" (blog), Palantir, Medium, March 17, 2025, <https://blog.palantir.com/palantirs-recommendations-to-the-white-house-ostp-on-developing-an-ai-action-plan-80d2f1821e8b>.
- Joel Kaplan, "More Speech and Fewer Mistakes," Meta, January 7, 2025, <https://about.fb.com/news/2025/01/meta-more-speech-fewer-mistakes/>.
- Paul Smith-Goodson and Matt Kimball, "The Stargate Project: Trump Touts \$500 Billion Bid for AI Dominance," *Forbes*, January 30, 2025, <https://www.forbes.com/sites/moorinsights/2025/01/30/the-stargate-project-trump-touts-500-billion-bid-for-ai-dominance/>.
- Elena Moore, Camila Domonoske, and Jeongyoon Han, "Trump Taps Musk to Lead a 'Department of Government Efficiency' with Ramaswamy," NPR, November 12, 2024, <https://www.npr.org/2024/11/12/g-s1-33972/trump-elon-musk-vivek-ramaswamy-doge-government-efficiency-deep-state>.

PAGE 46

- Trump, "Advancing Artificial Intelligence Education for American Youth."

PAGE 47

- U.S. Department of Labor, US Department of Labor releases AI literacy framework providing foundational content areas, delivery principles to guide nationwide efforts, February 13, 2026, <https://www.dol.gov/newsroom/releases/eta/eta20260213>.

- U.S. Department of Labor's AI Literacy Framework, February 2026, <https://www.dol.gov/sites/dolgov/files/OPA/newsreleases/2026/02/ETA-20260212-hi.jpg>.
- U.S. Department of Labor, US Department of Labor promotes AI literacy across the American workforce, August 26, 2025, <https://www.dol.gov/newsroom/releases/osec/osec20250826>.
- Advancing Artificial Intelligence in Education for American Youth, April 23, 2025, <https://www.whitehouse.gov/presidential-actions/2025/04/advancing-artificial-intelligence-education-for-american-youth/>.
- Dear Colleague Letter, Guidance on the Use of Federal Grant Funds to Improve Education Outcomes Using Artificial Intelligence, July 22, 2025, <https://www.ed.gov/media/document/oepd-ai-dear-colleague-letter-7222025-110427.pdf>.
- Presidential AI Challenge, <https://www.ai.gov/initiatives/presidential-challenge>.

PAGE 48

- S.146 — 119th Congress (2025-2026), <https://www.congress.gov/bill/119th-congress/senate-bill/146>.
- S.836 — 119th Congress (2025-2026), <https://www.congress.gov/bill/119th-congress/senate-bill/836>.
- H.R.7757 — 119th Congress (2025-2026), <https://www.congress.gov/bill/119th-congress/house-bill/7757/text>.
- S.1748 — 119th Congress (2025-2026), <https://www.congress.gov/bill/119th-congress/senate-bill/1748/text>.
- S.3062 — 119th Congress (2025-2026), <https://www.congress.gov/bill/119th-congress/senate-bill/3062/text>.
- S.2714 — 119th Congress (2025-2026), <https://www.congress.gov/bill/119th-congress/senate-bill/2714/text>.
- Ensuring a National Framework on Artificial Intelligence, <https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy/>.
- Promoting Advanced Artificial Intelligence Innovation and Security, Executive Order, June 2, 2026, <https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/>.

PAGE 49

- Multistate AI, “Artificial Intelligence (AI) Legislation Tracker 2026: All 50 States”, <https://www.multistate.ai/artificial-intelligence-ai-legislation>.
- Hillary Rinaldi, Kelly Branning, Annika Turnquist, Peter Cohen, Screen Time Legislation Is Moving Fast. Here’s What’s Actually Enacted., May 15, 2026, <https://whiteboardadvisors.com/screen-time-legislation-is-moving-fast-heres-whats-actually-enacted/>.
- State policies on cellphone use in K-12 public schools, last updated June 9, 2026, https://ballotpedia.org/State_policies_on_cellphone_use_in_K-12_public_schools.

PAGE 51

- “Engage With a Community of AI and Education Leaders,” TeachAI, <https://www.teachai.org/community>.
- “Foundational Policy Ideas for AI in Education,” TeachAI, <https://www.teachai.org/policy-ideas>.
- AI Policy Lab (blog), InnovateEDU, <https://www.aipolicylab.org/blog>.
- Erin Mote, interview by Bellwether, June 24, 2024.

PAGE 52

- ILO Group, *Framework for Implementing Artificial Intelligence (AI) in State Education Agencies (SEAs)*, 2024, https://www.ilogroup.com/wp-content/uploads/2024/06/Framework-for-Implementing-Artificial-Intelligence-AI-in-State-Education-Agencies-SEAs_v1.0.pdf.
- NCEE, *Framework for AI-Powered Learning Environments*, 2024, <https://ncee.org/whitepaper/framework-for-ai-powered-learning-environments/>.

PAGE 53

- National Education Association, *Report of the NEA Task Force on Artificial Intelligence in Education*, April 2024, https://www.nea.org/sites/default/files/2024-06/report_of_the_nea_task_force_on_artificial_intelligence_in_education_ra_2024.pdf.
- AFT, *Commonsense Guardrails for Using Advanced Technology in Schools*, 2024, https://www.aft.org/sites/default/files/media/documents/2024/Commonsense_Guardrails_AI_0604.pdf.

PAGE 54

- EDSAFE AI Alliance, *What Is the EDSAFE AI SAFE Framework?*, https://www.edsafeai.org/_files/ugd/5be6a9_b9a2f0289f5f475db2fe80d90f752f3f.pdf.
- Erin Mote, interview by Bellwether, June 24, 2024.

PAGE 55

- West Virginia Department of Education, *Guidance, Considerations, & Intentions for the Use of Artificial Intelligence in West Virginia Schools*, May 2024, <https://wvde.us/wp-content/uploads/2024/03/30438-WVDE-AI-Guidance-v1.pdf>.
- Santa Ana Unified School District, *SAUSD’s Artificial Intelligence Compass*, <https://www.sausd.us/cms/lib/CA01000471/Centricity/Domain/6/DRAFT%20-%20SAUSDs%20AI%20Compass%20Your%20Guide%20to%20Navigating%20an%20AI%20World%209.pdf>.

PAGE 56

- Commonwealth of Virginia, *Guidelines for AI Integration Throughout Education in the Commonwealth of Virginia*, 2024, <https://www.education.virginia.gov/media/governorviriniagov/secretary-of-education/pdf/AI-Education-Guidelines.pdf>.
- El Segundo Unified School District, *Board Policy 6163.4.1: Student Use of Emerging Technologies, Including Generative Artificial Intelligence*, <https://www.aipolicylab.org/post/model-board-policies-from-el-segundo-usd>.

PAGE 57

- Oregon Department of Education, *Generative Artificial Intelligence (AI) in K-12 Classrooms*, [https://www.oregon.gov/ode/educator-resources/teachingcontent/Documents/ODE_Generative_Artificial_Intelligence_\(AI\)_in_K-12_Classrooms_2023.pdf](https://www.oregon.gov/ode/educator-resources/teachingcontent/Documents/ODE_Generative_Artificial_Intelligence_(AI)_in_K-12_Classrooms_2023.pdf).
- “Artificial Intelligence - Principles and Beliefs,” Peninsula School District, <https://www.psd401.net/ai>.

PAGE 58

- El Segundo Unified School District, *Board Policy 6163.4.1*.

PAGE 59

- Cardona et al., *Artificial Intelligence*.
- Washington Office of Superintendent of Public Instruction, *Human-Centered AI Guidance for K–12 Public Schools*, March 27, 2024, <https://ospi.k12.wa.us/sites/default/files/2024-04/human-centered-ai-guidance-k-12-edition-2.pdf>.

PAGE 60

- Chicago Public Schools, *AI Guidebook*, https://static1.squarespace.com/static/64398599b0c21f1705fb8fb3/t/66bf92efe9cc2840f8a379d4/1723831025576/Chicago+Public+Schools_AI+Guidebook.pdf.

PAGE 61

- "The State of AI in Education 2025," Carnegie Learning, 2025, <https://discover.carnegielearning.com/hubfs/PDFs/Whitepaper%20and%20Guide%20PDFs/2025-AI-in-Ed-Report.pdf?hsLang=en>.
- RAND, *AI Use in Schools Is Quickly Increasing but Guidance Lags Behind*, September 2025, https://www.rand.org/pubs/research_reports/RR4180-1.html.
- "School Pulse Panel: Surveying High-Priority, Education-Related Topics," National Center for Education Statistics, <https://nces.ed.gov/surveys/spp/results.asp#technology-dec24-chart-4>.
- Lauraine Langreo, "'We're at a Disadvantage,' and Other Teacher Sentiments on AI," *Education Week*, October 29, 2024, <https://www.edweek.org/technology/were-at-a-disadvantage-and-other-teacher-sentiments-on-ai/2024/10>.

PAGE 62

- Cardona et al., *Artificial Intelligence*.
- TeachAI, *AI Guidance for Schools Toolkit*, <https://www.teachai.org/toolkit-principles>.
- Educating All Learners Alliance and New America, *Prioritizing Students With Disabilities in AI Policy*, <https://drive.google.com/file/d/1iaY6s466mlvzo-9SmcuKbohxF1Pc274/view>.
- Wayne Holmes, "The Unintended Consequences of Artificial Intelligence and Education," *Education International Research*, updated March 22, 2024, <https://www.ei-ie.org/en/item/28115:the-unintended-consequences-of-artificial-intelligence-and-education>.

PAGE 63

- All4Ed, *Demystifying Artificial Intelligence (AI) for K-12*, <https://all4ed.org/future-ready-schools/emerging-practices-guides/demystifying-artificial-intelligence-ai-for-k-12/#challenges-and-opportunities>.
- Michael Klein and Zac Chase, *K-12 Digital Infrastructure Brief: Defensible and Resilient*, U.S. Department of Education, 2023, https://tech.ed.gov/files/2023/08/DOEd-Report_20230804_-508c.pdf.

- Rebecca Torchia, "TCEA 2024: Planning and Administrator Support Are Necessary to Sustain Devices," *EdTech*, February 7, 2024, <https://edtechmagazine.com/k12/article/2024/02/tcea-2024-planning-and-administrator-support-are-necessary-sustain-devices>.
- Shawn Augenstein, "The Stumbling Blocks for Artificial Intelligence in K-12 Education," *EdTech*, May 7, 2024, <https://edtechmagazine.com/k12/article/2024/05/stumbling-blocks-artificial-intelligence-k-12-education>.
- Naaz Modan, "Supreme Court to Hear FCC E-Rate Case on School Reimbursements," *K-12 Dive*, June 18, 2024, <https://www.k12dive.com/news/supreme-court-fcc-e-rate-school-reimbursements/719257/>.
- Holmes, "Unintended Consequences."

PAGE 65

- "The State of AI in Education 2025," Carnegie Learning.
- Kaufman, Woo, Eagan, Lee, and Kassan, "Uneven Adoption of Artificial Intelligence Tools Among U.S. Teachers and Principals in the 2023-2024 School Year."
- Langreo, "'We're at a Disadvantage,' and Other Teacher Sentiments on AI."
- Christopher Joseph Doss et al, *AI Use in Schools is Quickly Increasing but Guidance Lags Behind*, RAND, Sept 30, 2025, https://www.rand.org/pubs/research_reports/RR4180-1.html.

PAGE 66

- "The State of AI in Education 2025," Carnegie Learning.
- Kaufman, Woo, Eagan, Lee, and Kassan, "Uneven Adoption of Artificial Intelligence Tools Among U.S. Teachers and Principals in the 2023-2024 School Year."
- Langreo, "'We're at a Disadvantage,' and Other Teacher Sentiments on AI."
- Caitlynn Peetz, "What Teacher PD on AI Should Look Like. Some Early Models Are Emerging," *Education Week*, December 9, 2024, <https://www.edweek.org/technology/what-teacher-pd-on-ai-should-look-like-some-early-models-are-emerging/2024/12>.

PAGE 67

- Langreo, "'We're at a Disadvantage,' and Other Teacher Sentiments on AI."
- "The State of AI in Education 2025," Carnegie Learning.

- Kaufman, Woo, Eagan, Lee, and Kassan, “Uneven Adoption of Artificial Intelligence Tools Among U.S. Teachers and Principals in the 2023–2024 School Year.”
- “School Pulse Panel: Surveying High-Priority, Education-Related Topics,” National Center for Education Statistics.
- Peetz, “What Teacher PD on AI Should Look Like. Some Early Models Are Emerging.”
- Melissa Kay Diliberti, Robin J. Lake, and Steven R. Weiner, “More Districts are Training Teachers on Artificial Intelligence,” RAND, April 8, 2025, https://www.rand.org/pubs/research_reports/RRA956-31.html.

PAGE 68

- Diliberti et al., *Using Artificial Intelligence Tools*, 10.

PAGE 70

- AI for Equity, *AI Exemplar: Washington Leadership Academy*, https://docs.google.com/document/d/1eRznu81AlrAILKtQc58g2wTO2m5_-qtPXy306QzQ2Qg/edit.

PAGE 73

- “Maximizing the Value of Your Tech,” Leading Educators, <https://www.valueedtech.org/the-framework>.
- Opportunity Labs and F3 Law, *Procurement Benchmarks for AI in K-12 Education: Using Educational Purchasing Power to Shift Incentives Towards Technology Development for Child Wellbeing*, 2024, <https://static1.squarespace.com/static/6596cbea8bbae0486d04d7b9/t/6675c2f1e031306bd6ef35c6/1718993650308/K-12+Procurement+Benchmarks+June+2024.pdf>.

PAGE 74

- AI Use by Tweens and Teens, Common Sense Media, 2026, <https://www.commonsensemedia.org/sites/default/files/research/report/2026-ai-use-by-tweens-and-teens-1.pdf>.
- Michelle Faverio and Olivia Sidoti, *Teens, Social Media and AI Chatbots 2025*, Pew Research Center, December 9, 2025, <https://www.pewresearch.org/internet/2025/12/09/teens-social-media-and-ai-chatbots-2025/#use-of-ai-chatbots>.

PAGE 75

- “Research Brief: Teens, Trust, and Technology in the Age of AI,” Common Sense Media, January 29, 2025, <https://www.commonsensemedia.org/research/research-brief-teens-trust-and-technology-in-the-age-of-ai>.
- “Voices of Gen Z: Preparing the Heartland for an AI Future,” Gallup and Walton Family Foundation, 2025, <https://www.gallup.com/analytics/651674/gen-z-research.aspx>.
- The AI Paradox: More Exposure, Less Confidence Among Gen Z, 2026, Gallup, Walton Family Foundation, GSV, <https://www.gallup.com/analytics/651674/gen-z-research.aspx>.

PAGE 76

- “College Students and ChatGPT Adoption in the US,” OpenAI, February 20, 2025, <https://openai.com/global-affairs/college-students-and-chatgpt/>.
- “Anthropic Education Report: How University Students Use Claude,” Anthropic, April 8, 2025, <https://www.anthropic.com/news/anthropic-education-report-how-university-students-use-claude>.
- Kylie Robison, “OpenAI and Anthropic Are Fighting Over College Students With Free AI,” The Verge, April 3, 2025, <https://www.theverge.com/ai-artificial-intelligence/641193/openai-anthropic-education-tool-college>.

PAGE 77

- “AI Initiative: Understanding the Impact of AI on Our Kids,” Common Sense Media, <https://www.commonsensemedia.org/ai>.

PAGE 78

- Monica P. Bhatt et al., “Can Technology Facilitate Scale? Evidence From a Randomized Evaluation of High Dosage Tutoring,” NBER Working Paper Series 32510, National Bureau of Economic Research, May 2024, <https://www.nber.org/papers/w32510>.
- Hamsa Bastani et al., “Generative AI Can Harm Learning,” SSRN, July 15, 2024, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4895486.
- Owen Henkel et al., “Can Large Language Models Make the Grade? An Empirical Study Evaluating LLMs Ability to Mark Short Answer Questions in K-12 Education,” in *Proceedings of the Eleventh ACM Conference on Learning @ Scale (L@S '24)* (New York, NY: Association for Computing Machinery, 2024), 300–304, <https://doi.org/10.1145/3657604.3664693>.
- Jake Bryant et al. “How Artificial Intelligence Will Impact K–12 Teachers,” McKinsey & Company, January 14, 2020, <https://www.mckinsey.com/industries/education/our-insights/how-artificial-intelligence-will-impact-k-12-teachers>.

PAGE 79

- National Academies of Sciences, Engineering, and Medicine, *How People Learn II: Learners, Contexts, and Cultures* (Washington, DC: The National Academies Press, 2018), 109, 126, 156, 117, 155, <https://doi.org/10.17226/24783>.

PAGE 80

- Michael B. Robb and Supreet Mann, *Talk, Trust, and Trade-Offs: How and Why Teens Use AI Companions* (San Francisco, CA: Common Sense Media, 2025), https://www.common Sense Media.org/sites/default/files/research/report/talk-trust-and-trade-offs_2025_web.pdf.
- Scott Rosenberg, "Parents Sue OpenAI Over Teen's Suicide," *Axios*, August 26, 2025, <https://www.axios.com/2025/08/26/parents-sue-openai-chatgpt>.
- Hadas Gold, "More Families Sue Character.AI Developer, Alleging App Played a Role in Teens' Suicide and Suicide Attempt," *CNN*, September 16, 2025, <https://www.cnn.com/2025/09/16/tech/character-ai-developer-lawsuit-teens-suicide-and-suicide-attempt>.
- Rihitu Chatterjee, "Their Teenage Sons Died by Suicide. Now, They Are Sounding an Alarm About AI Chatbots," *NPR*, September 19, 2025, <https://www.npr.org/sections/shots-health-news/2025/09/19/nx-s1-5545749/ai-chatbots-safety-openai-meta-characterai-teens-suicide>.
- Jeff Horwitz, "Meta's AI Rules Have Let Bots Hold 'Sensual' Chats With Kids, Offer False Medical Info," *Reuters*, August 14, 2025, <https://www.reuters.com/investigates/special-report/meta-ai-chatbot-guidelines/>.
- "FTC Launches Inquiry Into AI Chatbots Acting as Companions," news release, Federal Trade Commission, September 11, 2025, <https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-launches-inquiry-ai-chatbots-acting-companions>.
- "Bipartisan Coalition of State Attorneys General Issues Letter to AI Industry Leaders on Child Safety," news release, National Association of Attorneys General, August 26, 2025, <https://www.naag.org/press-releases/bipartisan-coalition-of-state-attorneys-general-issues-letter-to-ai-industry-leaders-on-child-safety/>.
- Jamie Bernardi, "Friends for Sale: The Rise and Risks of AI Companions," *Ada Lovelace Institute*, January 23, 2025, <https://www.adalovelaceinstitute.org/blog/ai-companions/>.

- John Sanford, "Why AI Companions and Young People Can Make for a Dangerous Mix," *Stanford Report*, August 27, 2025, <https://news.stanford.edu/stories/2025/08/ai-companions-chatbots-teens-young-people-risks-dangers-study>.
- Justine Gluck and Rafal Fryc, 2026 Chatbot Legislation Tracker, <https://fpf.org/2026-chatbot-legislation-tracker/>.

PAGE 81

- Laurence Holt, "The 5 Percent Problem," *Education Next*, April 11, 2024, <https://www.educationnext.org/5-percent-problem-online-mathematics-programs-may-benefit-most-kids-who-need-it-least/>.

PAGE 82

- Emma Kate Fittes, "LAUSD and AllHere: 4 Takeaways Amid New Doubts About the Far-Reaching AI Project," *EdWeek Market Brief*, July 2, 2024, <https://marketbrief.edweek.org/strategy-operations/laUSD-and-allhere-4-takeaways-amid-new-doubts-about-the-far-reaching-ai-project/2024/07>.
- Ben Chapman, "Turmoil Surrounds LA's New AI Student Chatbot as Tech Firm Furloughs Staff Just 3 Months After Launch," *The 74*, June 26, 2024, <https://www.the74million.org/article/turmoil-surrounds-las-new-ai-student-chatbot-as-tech-firm-furloughs-staff-just-3-months-after-launch/>.
- Lauraine Langreo, "Los Angeles Unified Bets Big on 'Ed,' an AI Tool for Students," *Education Week*, March 21, 2024, <https://www.edweek.org/technology/los-angeles-unified-bets-big-on-ed-an-ai-tool-for-students/2024/03>.
- Howard Blume, "LAUSD Shuttles New AI Chatbot as Developer Goes Out of Business," *Government Technology*, July 3, 2024, <https://www.govtech.com/education/k-12/laUSD-shuttles-new-ai-chatbot-when-developer-goes-out-of-business>.

PAGE 83

- Chapman, "New AI Student Chatbot."
- Fittes, "LAUSD and AllHere."
- Mark Keierleber, "L.A. Schools Probe Charges Its Hyped, Now-Defunct AI Chatbot Misused Student Data," *The 74*, July 10, 2024, <https://www.the74million.org/article/chatbot-los-angeles-whistleblower-allhere-ai/>.



Contents

Introduction

Key Takeaways

Overview

Platform

Policy

User Experience

Conclusion

Sources

Acknowledgments

About the Authors

About Bellwether

Interviews

Original interviews were conducted May through August 2024. The following individuals lent their perspectives and expertise on this series, including:

Aaron Cuny
AI for Equity

Aasim Shabazz
Twin Cities Innovation Alliance

Adam Feiler
DREAM PCS

Adeel Khan
MagicSchool AI

Alejandro Gibes de Gac
Paloma

**Amanda Lenhart, Beth Sears,
Ellen Pack, Merve Lapus,
Samantha Ketcham, and
Yvette Renteria**
Common Sense Media

**Angela Duckworth, Ethan Mollick,
and Lilach Mollick**
University of Pennsylvania

Benjamin Riley
Cognitive Resonance

Bobby Moore and Luis Pérez
CAST

Bree Dusseault and Robin Lake
Center on Reinventing Public Education

Brent Maddin
ASU School of Ed

Brent Milne and Tom Fischhaber
Saga Education

Cameron White and Erin Stark
NewSchools

Claire Zau
GSV Ventures

**Dale Schmidt, Elaine Perea,
Robert Runcie, and Shayla Cannady**
Chiefs for Change

Dan Effland
Summit PCS

Dan Meyer
Amplify and Mathworlds

David Fields
Northeastern University

Dylan Kane
Five Twelve Thirteen Substack

Erin Mote
InnovateEDU

Jack Shaw
Comprendo.dev

Jennie Magiera and Kristal Ayres
Google, Education Impact

Jeremy Roschelle
Digital Promise

Jimmy Fisher
Seckinger High School

Josh Clark
Landmark School

Kim Smith
LearnerStudio

Kristen DiCerbo
Khan Academy

Kumar Garg
Renaissance Philanthropy

Lorne Rodriguez and Peter Leonard
Chicago Public Schools

Maheen Sahoo and Peter Gault
Quill

Matthew Marino
University of Central Florida

Michelle Rojas-Soto
KIPP SoCal

Pat Yongpradit
Code.org

Rebecca Finlay
Partnership on AI

Rob Waldron
Curriculum Associates

Stacy Kane
Washington Leadership Academy

Sue Krause
Cognia

Yusuf Ahmad
Playlab.ai

About the Authors



AMY CHEN KULESA

Amy Chen Kulesa is a senior adviser at Bellwether. She can be reached at amy@luminosityadvising.com.



MICHELLE CROFT

Michelle Croft is an associate partner at Bellwether. She can be reached at michelle.croft@bellwether.org.



MARISA MISSION

Marisa Mission is a senior analyst at Bellwether. She can be reached at marisa.mission@bellwether.org.



BRIAN ROBINSON

Brian Robinson is an associate partner at Bellwether. He can be reached at brian.robinson@bellwether.org.



LINEA HARDING

Linea Harding is a senior analyst at Bellwether. She can be reached at linea.harding@bellwether.org.



MARY K. WELLS

Mary K. Wells is a co-founder and senior partner at Bellwether and leads the organization's work on nontraditional and out-of-school learning. She can be reached at mary@bellwether.org.



ANDREW J. ROTHERHAM

Andrew J. Rotherham is a co-founder and senior partner at Bellwether. He can be reached at andy@bellwether.org.



JOHN BAILEY

John Bailey is a nonresident senior fellow at American Enterprise Institute. He can be reached at john@vestigopartners.com.

About Bellwether

Bellwether is a national nonprofit that works to transform education to ensure young people — especially those furthest from opportunity — achieve outcomes that lead to fulfilling lives and flourishing communities. Founded in 2010, we help mission-driven partners accelerate their impact, inform and influence policy and program design, and bring leaders together to drive change on education's most pressing challenges. For more, visit bellwether.org.

ACKNOWLEDGMENTS

We would like to thank the many individuals who gave their time and shared their knowledge with us to inform our work, including interview participants for being generous in providing their perspective and expertise. Thank you to the Bezos Family Foundation, Chan Zuckerberg Initiative, Charles and Helen Schwab Foundation, the Charter School Growth Fund, Oak Foundation, Overdeck Family Foundation, and Raikes Foundation for their financial support of the September 2024, May 2025, and January 2026 editions of this project. Thank you also to the Charter School Growth Fund and the Walton Family Foundation for their financial support of the August 2026 edition of this project.

We would also like to thank our Bellwether colleagues Alex Spurrier, Daniela Torre Gibney, Christine Wade, Lynne Wells Graziano, Mark Baxter, Priyanka Patel, and Alexis Richardson for their input, and our former colleagues Shruti Nagarajan, Zoe Cuddy, Mandy Berman, Kaitlin Beeson, and Janine Sandy for their support. Thank you to Amy Ribock, Kate Stein, Andy Jacob, McKenzie Maxson, Esta Sherr, Temim Fruchter, Julie Nguyen, and Amber Walker for shepherding and disseminating this work, and to Super Copy Editors.

The contributions of these individuals and entities significantly enhanced our work; however, any errors in fact or analysis remain the responsibility of the authors.

DISCLOSURE

Bellwether works with organizations and leaders who share our viewpoint-diverse commitment to improving education and advancing equity for all young people — regardless of identity, circumstance, or background. As part of our commitment to transparency, a list of Bellwether clients and funders since our founding in 2010 is publicly available on our website. An organization's name appearing on our list of clients and funders does not imply any endorsement of or by Bellwether.

Similarly, examples of tools, companies, and nonprofits featured in this three-part AI analysis do not imply any endorsement of or by Bellwether for specific products or services.

Separate from the creation of this three-part AI analysis, Bellwether has worked with CRPE and receives funding from the Walton Family Foundation. In addition, a co-founder and partner is an adviser to RAND.



© 2026 Bellwether

- © This report carries a Creative Commons license, which permits noncommercial reuse of content when proper attribution is provided. This means you are free to copy, display, and distribute this work, or include content from this report in derivative works, under the following conditions:
- ① **Attribution.** You must clearly attribute the work to Bellwether and provide a link back to the publication at www.bellwether.org.
 - ② **Noncommercial.** You may not use this work for commercial purposes without explicit prior permission from Bellwether.
 - ③ **Share Alike.** If you alter, transform, or build upon this work, you may distribute the resulting work only under a license identical to this one.

For the full legal code of this Creative Commons license, please visit www.creativecommons.org. If you have any questions about citing or reusing Bellwether content, please contact us.