



OFFICE OF
Educational Technology

Artificial Intelligence and the Future of Teaching and Learning

Insights and Recommendations

May 2023



Artificial Intelligence and the Future of Teaching and Learning

Miguel A. Cardona, Ed.D.

Secretary, U.S. Department of Education

Roberto J. Rodríguez

Assistant Secretary, Office of Planning, Evaluation, and Policy Development

Kristina Ishmael

Deputy Director, Office of Educational Technology

May 2023

Examples Are Not Endorsements

This document contains examples and resource materials that are provided for the user's convenience. The inclusion of any material is not intended to reflect its importance nor is it intended to endorse any views expressed or products or services offered. These materials may contain the views and recommendations of various subject matter experts as well as hypertext links, contact addresses, and websites to information created and maintained by other public and private organizations. The opinions expressed in any of these materials do not necessarily reflect the positions or policies of the U.S. Department of Education. The U.S. Department of Education does not control or guarantee the accuracy, relevance, timeliness, or completeness of any information from other sources that are included in these materials. Other than statutory and regulatory requirements included in the document, the contents of this guidance do not have the force and effect of law and are not meant to bind the public.

Contracts and Procurement

This document is not intended to provide legal advice or approval of any potential federal contractor's business decision or strategy in relation to any current or future federal procurement and/or contract. Further, this document is not an invitation for bid, request for proposal, or other solicitation.

Licensing and Availability

This report is in the public domain and available on the U.S. Department of Education's (Department's) website at <https://tech.ed.gov>.

Requests for alternate format documents such as Braille or large print should be submitted to the Alternate Format Center by calling 1-202-260-0852 or by contacting the 504 coordinator via [email](mailto:om_eeos@ed.gov) at om_eeos@ed.gov.

Notice to Limited English Proficient Persons

If you have difficulty understanding English, you may request language assistance services for Department information that is available to the public. These language assistance services are available free of charge. If you need more information about interpretation or translation services, please call 1-800-USA-LEARN (1-800-872-5327) (TTY: 1-800-437-0833); email us at Ed.Language.Assistance@ed.gov; or write to U.S. Department of Education, Information Resource Center, LBJ Education Building, 400 Maryland Ave. SW, Washington, DC 20202.

How to Cite

While permission to reprint this publication is not necessary, the suggested citation is as follows:

U.S. Department of Education, Office of Educational Technology, *Artificial Intelligence and Future of Teaching and Learning: Insights and Recommendations*, Washington, DC, 2023.

This report is available at <https://tech.ed.gov>

Table of Contents

Introduction.....	1
Rising Interest in AI in Education	1
Three Reasons to Address AI in Education Now.....	2
Toward Policies for AI in Education	3
Building Ethical, Equitable Policies Together.....	6
Guiding Questions.....	6
Foundation 1: Center People (Parents, Educators, and Students).....	6
Foundation 2: Advance Equity.....	7
Foundation 3: Ensure Safety, Ethics, and Effectiveness	8
Foundation 4: Promote Transparency	9
Overview of Document	10
What is AI?	11
Perspective: Human-Like Reasoning.....	12
Perspective: An Algorithm that Pursues a Goal.....	12
Perspective: Intelligence Augmentation	14
Definition of "Model"	14
Insight: AI Systems Enable New Forms of Interaction.....	15
Key Recommendation: Human in the Loop AI.....	16
Learning.....	18
Insight: AI Enables Adaptivity in Learning.....	18
Intelligent Tutoring Systems: An Example of AI Models	19
Important Directions for Expanding AI-Based Adaptivity.....	20
A Duality: Learning With and About AI.....	22
A Challenge: Systems Thinking About AI in Education.....	22
Open Questions About AI for Learning.....	23
Key Recommendation: Seek AI Models Aligned to a Vision for Learning.....	24
Teaching.....	25
Always Center Educators in Instructional Loops	25
Insight: Using AI to Improve Teaching Jobs	26
Preparing and Supporting Teachers in Planning and Reflecting.....	29
Designing, Selecting, and Evaluating AI Tools.....	30
Challenge: Balancing Human and Computer Decision-Making	30
Challenge: Making Teaching Jobs Easier While Avoiding Surveillance	31
Challenge: Responding to Students' Strengths While Protecting Their Privacy	32
Questions Worth Asking About AI for Teaching.....	34
Key Recommendation: Inspectable, Explainable, Overridable AI.....	34

Formative Assessment	37
Building on Best Practices	37
Implications for Teaching and Learning	38
Insight: AI Can Enhance Feedback Loops	39
An Example: Automated Essay Scoring	40
Key Opportunities for AI in Formative Assessment	41
Key Recommendation: Harness Assessment Expertise to Reduce Bias	42
Related Questions	43
Research and Development	44
Insight: Research Can Strengthen the Role of Context in AI	44
Attention to the Long Tail of Learner Variability	46
Partnership in Design-Based Research	47
Re-thinking Teacher Professional Development	48
Connecting with Public Policy	49
Key Recommendation: Focus R&D on Addressing Context	50
Ongoing Questions for Researchers	50
Desired National R&D Objectives	51
Recommendations	52
Insight: Aligning AI to Policy Objectives	52
Calling Education Leaders to Action	53
Recommendation #1: Emphasize Humans in the Loop	53
Recommendation #2: Align AI Models to a Shared Vision for Education	54
Recommendation #3: Design Using Modern Learning Principles	56
Recommendation #4: Prioritize Strengthening Trust	57
Recommendation #5: Inform and Involve Educators	57
Recommendation #6: Focus R&D on Addressing Context and Enhancing Trust and Safety	59
Recommendation #7: Develop Education-Specific Guidelines and Guardrails	60
Next Steps	60
Common Acronyms and Abbreviations	62
Acknowledgements	63
References	64

Introduction

The U.S. Department of Education (Department) is committed to supporting the use of technology to improve teaching and learning and to support innovation throughout educational systems. This report addresses the clear need for sharing knowledge and developing policies for “Artificial Intelligence,” a rapidly advancing class of foundational capabilities which are increasingly embedded in all types of educational technology systems and are also available to the public. We will consider “educational technology” (edtech) to include both (a) technologies specifically designed for educational use, as well as (b) general technologies that are widely used in educational settings. Recommendations in this report seek to engage teachers, educational leaders, policy makers, researchers, and educational technology innovators and providers as they work together on pressing policy issues that arise as Artificial Intelligence (AI) is used in education.

AI can be defined as “automation based on associations.” When computers automate reasoning based on associations in data (or associations deduced from expert knowledge), two shifts fundamental to AI occur and shift computing beyond conventional edtech: (1) from capturing data to *detecting patterns* in data and (2) from providing access to instructional resources to *automating decisions* about instruction and other educational processes. Detecting patterns and automating decisions are leaps in the level of responsibilities that can be delegated to a computer system. The process of developing an AI system may lead to bias in how patterns are detected and unfairness in how decisions are automated. Thus, educational systems must govern their use of AI systems. This report describes opportunities for using AI to improve education, recognizes challenges that will arise, and develops recommendations to guide further policy development.

Rising Interest in AI in Education

Today, many priorities for improvements to teaching and learning are unmet. Educators seek technology-enhanced approaches addressing these priorities that would be safe, effective, and scalable. Naturally, educators wonder if the rapid advances in technology in everyday lives could help. Like all of us, educators use AI-powered services in their everyday lives, such as voice assistants in their homes; tools that can correct grammar, complete sentences, and write essays; and automated trip planning on their phones. Many educators are actively exploring AI tools as they are newly released to the public¹. Educators see opportunities to use AI-powered capabilities like speech recognition to increase the support available to students with disabilities, multilingual learners, and others who could benefit from greater adaptivity and personalization in digital tools for learning. They are exploring how AI can enable writing or improving lessons, as well as their process for finding, choosing, and adapting material for use in their lessons.

Educators are also aware of new risks. Useful, powerful functionality can also be accompanied with new data privacy and security risks. Educators recognize that AI can automatically produce output that is inappropriate or wrong. They are wary that the associations or automations created by AI may amplify unwanted biases. They have noted new ways in which students may

¹ Walton Family Foundation (March 1, 2023). Teachers and students embrace ChatGPT for education. <https://www.waltonfamilyfoundation.org/learning/teachers-and-students-embrace-chatgpt-for-education>

represent others' work as their own. They are well-aware of "teachable moments" and pedagogical strategies that a human teacher can address but are undetected or misunderstood by AI models. They worry whether recommendations suggested by an algorithm would be fair. Educators' concerns are manifold. Everyone in education has a responsibility to harness the good to serve educational priorities while also protecting against the dangers that may arise as a result of AI being integrated in edtech.

To develop guidance for edtech, the Department works closely with educational constituents. These constituents include educational leaders—teachers, faculty, support staff, and other educators—researchers; policymakers; advocates and funders; technology developers; community members and organizations; and, above all, learners and their families/caregivers. Recently, through its activities with constituents, the Department noticed a sharp rise in interest and concern about AI. For example, a 2021 field scan found that developers of all kinds of technology systems—for student information, classroom instruction, school logistics, parent-teacher communication, and more—expect to add AI capabilities to their systems. Through a series of four listening sessions conducted in June and August 2022 and attended by more than 700 attendees, it became clear that constituents believe that action is required now in order to get ahead of the expected increase of AI in education technology—and they want to roll up their sleeves and start working together. In late 2022 and early 2023, the public became aware of new generative AI chatbots and began to explore how AI could be used to write essays, create lesson plans, produce images, create personalized assignments for students, and more. From public expression in social media, at conferences, and in news media, the Department learned more about risks and benefits of AI-enabled chatbots. And yet this report will not focus on a specific AI tool, service, or announcement, because AI-enabled systems evolve rapidly. Finally, the Department engaged the educational policy expertise available internally and in its relationships with AI policy experts to shape the findings and recommendations in this report.

Three Reasons to Address AI in Education Now

"I strongly believe in the need for stakeholders to understand the cyclical effects of AI and education. By understanding how different activities accrue, we have the ability to support virtuous cycles. Otherwise, we will likely allow vicious cycles to perpetuate."

—Lydia Liu

During the listening sessions, constituents articulated three reasons to address AI now:

First, AI may enable achieving educational priorities in better ways, at scale, and with lower costs. Addressing varied unfinished learning of students due to the pandemic is a policy priority, and AI may improve the adaptivity of learning resources to students' strengths and needs. Improving teaching jobs is a priority, and via automated assistants or other tools, AI may provide teachers greater support. AI may also enable teachers to extend the support they offer to individual students when they run out of time. Developing resources that are responsive to the knowledge and experiences students bring to their learning—their community and cultural assets—is a priority, and AI may enable greater customizability of curricular resources to meet local needs.

As seen in voice assistants, mapping tools, shopping recommendations, essay-writing capabilities, and other familiar applications, AI may enhance educational services.

Second, urgency and importance arise through awareness of system-level risks and anxiety about potential future risks. For example, students may become subject to greater surveillance. Some teachers worry that they may be replaced—to the contrary, the Department firmly rejects the idea that AI could replace teachers. Examples of discrimination from algorithmic bias are on the public's mind, such as a voice recognition system that doesn't work as well with regional dialects, or an exam monitoring system that may unfairly identify some groups of students for disciplinary action. Some uses of AI may be infrastructural and invisible, which creates concerns about transparency and trust. AI often arrives in new applications with the aura of magic, but educators and procurement policies require that edtech show efficacy. AI may provide information that appears authentic, but actually is inaccurate or lacking a basis in reality. Of the highest importance, AI brings new risks in addition to the well-known data privacy and data security risks, such as the risk of scaling pattern detectors and automations that result in "algorithmic discrimination" (e.g., systematic unfairness in the learning opportunities or resources recommended to some populations of students).

Third, urgency arises because of the scale of possible unintended or unexpected consequences. When AI enables instructional decisions to be automated at scale, educators may discover unwanted consequences. In a simple example, if AI adapts by speeding curricular pace for some students and by slowing the pace for other students (based on incomplete data, poor theories, or biased assumptions about learning), achievement gaps could widen. In some cases, the quality of available data may produce unexpected results. For example, an AI-enabled teacher hiring system might be assumed to be more objective than human-based résumé scoring. Yet, if the AI system relies on poor quality historical data, it might de-prioritize candidates who could bring both diversity and talent to a school's teaching workforce.

In summary, it is imperative to address AI in education *now* to realize key opportunities, prevent and mitigate emergent risks, and tackle unintended consequences.

Toward Policies for AI in Education

The 2023 [AI Index Report](#) from the Stanford Institute for Human-Centered AI has documented notable acceleration of investment in AI as well as an increase of research on ethics, including issues of fairness and transparency.² Of course, research on topics like ethics is increasing because problems are observed. Ethical problems will occur in education, too.³ The report found a striking interest in 25 countries in the number of legislative proposals that specifically include AI. In the United States, multiple executive orders are focused on ensuring AI is trustworthy and equitable, and the White House Office of Science and Technology Policy has introduced a

² Maslej, N., Fattorini, L., Brynjolfsson E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Ngo, H., Niebles, J.C., Parli, V., Shoham, Y., Wald, R., Clark, J. and Perrault, R., (2023). *The AI index 2023 annual report*. Stanford University: AI Index Steering Committee, Institute for Human-Centered AI.

³ Holmes, W. & Porayska-Pomsta, K. (Eds.) (2022). *The ethics of artificial intelligence in education*. Routledge. ISBN 978-0367349721


Blueprint for an AI Bill of Rights (Blueprint)⁴ that provides principles and practices that help achieve this goal. These initiatives, along with other AI-related policy activities occurring in both the executive and legislative branches, will guide the use of AI throughout all sectors of society. In Europe, the European Commission recently released Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators.⁵

AI is moving fast and heralding societal changes that require a national policy response. In addition to broad policies for all sectors of society, education-specific policies are needed to address new opportunities and challenges within existing frameworks that take into consideration federal student privacy laws (such as the Family Educational Rights and Privacy Act, or FERPA), as well as similar state related laws. AI also makes recommendations and takes actions automatically in support of student learning, and thus educators will need to consider how such recommendations and actions can comply with laws such as the Individuals with Disabilities Education Act (IDEA). We discuss specific policies in the concluding section.

Figure 1: Research about AI is growing rapidly. Other indicators, such as dollars invested and number of people employed, show similar trends.

Number of AI Publications by Field of Study (Excluding Other AI), 2010–21

Source: Center for Security and Emerging Technology, 2022 | Chart: 2023 AI Index Report



AI is advancing exponentially (see Figure 1), with powerful new AI features for generating images and text becoming available to the public, and leading to changes in how people create text and

⁴ White House Office of Science and Technology Policy (October 2022), *Blueprint for an AI bill of rights: Making automated systems work for the American people*. The White House Office of Science and Technology Policy. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>

⁵ European Commission, Directorate-General for Education, Youth, Sport and Culture. (2022). *Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators*, Publications Office of the European Union. <https://data.europa.eu/doi/10.2766/153756>

images⁶. The advances in AI are not only happening in research labs but also are making news in mainstream media and in educational-specific publications.

Researchers have articulated a range of concepts and frameworks for ethical AI⁷, as well as for related concepts such as equitable, responsible, and human-centered AI. Listening session participants called for building on these concepts and frameworks but also recognized the need to do more; participants noted a pressing need for guardrails and guidelines that make educational use of AI advances safe, especially given this accelerating pace of incorporation of AI into mainstream technologies. As policy development takes time, policy makers and educational constituents together need to start now to specify the requirements, disclosures, regulations, and other structures that can shape a positive and safe future for all constituents—especially students and teachers.

Policies are urgently needed to implement the following:

1. leverage automation to advance learning outcomes while protecting human decision making and judgment;
2. interrogate the underlying data quality in AI models to ensure fair and unbiased pattern recognition and decision making in educational applications, based on accurate information appropriate to the pedagogical situation;
3. enable examination of how particular AI technologies, as part of larger edtech or educational systems, may increase or undermine equity for students; and
4. take steps to safeguard and advance equity, including providing for human checks and balances and limiting any AI systems and tools that undermine equity.

⁶ Sharples, M. & Pérez y Pérez, R. (2022). *Story machines: How computers have become creative writers*. Routledge. ISBN 9780367751951

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

Building Ethical, Equitable Policies Together

In this report, we aim to build on the listening sessions the Department hosted to engage and inform all constituents involved in making educational decisions so they can prepare for and make better decisions about the role of AI in teaching and learning. AI is a complex and broad topic, and we are not able to cover everything nor resolve issues that still require more constituent engagement. This report is intended to be a starting point.

The opportunities and issues of AI in education are equally important in K-12, higher education, and workforce learning. Due to scope limitations, the examples in this report will focus on K-12 education. The implications are similar at all levels of education, and the Department intends further activities in 2023 to engage constituents beyond K-12 schools.

Guiding Questions

Understanding that AI increases automation and allows machines to do some tasks that only people did in the past leads us to a pair of bold, overarching questions:

1. What is our collective vision of a desirable and achievable educational system that leverages automation to advance learning while protecting and centering human agency?
2. How and on what timeline will we be ready with necessary guidelines and guardrails, as well as convincing evidence of positive impacts, so that constituents can ethically and equitably implement this vision widely?

In the Learning, Teaching, and Assessment sections of this report, we elaborate on elements of an educational vision grounded in what today's learners, teachers, and educational systems need, and we describe key insights and next steps required. Below, we articulate four key foundations for framing these themes. These foundations arise from what we know about the effective use of educational technology to improve opportunity, equity, and outcomes for students and also relate to the new Blueprint.

Foundation 1: Center People (Parents, Educators, and Students)

Education-focused AI policies at the federal, state, and district levels will be needed to guide and empower local and individual decisions about which technologies to adopt and use in schools and classrooms. Consider what is happening in everyday lives. Many of us use AI-enabled products because they are often better and more convenient. For example, few people want to use paper maps anymore; people find that technology helps us plan the best route to a destination more efficiently and conveniently. And yet, people often do not realize how much privacy they are giving up when they accept AI-enabled systems into their lives. AI will bring privacy and other risks that are hard to address only via individual decision making; additional protections will be needed.

There should be clear limits on the ability to collect, use, transfer, and maintain our personal data, including limits on targeted advertising. These limits should put the burden on platforms to minimize how much information they collect, rather than burdening Americans with reading fine print.⁸

As protections are developed, we recommend that policies center people, not machines. To this end, a first recommendation in this document (in the next section) is an emphasis on **AI with humans in the loop**. Teachers, learners, and others need to retain their agency to decide what patterns mean and to choose courses of action. The idea of humans in the loop builds on the concept of “Human Alternatives, Consideration, and Fallback” in the Blueprint and ethical concepts used more broadly in evaluating AI, such as preserving human dignity. A top policy priority must be establishing human in the loop as a requirement in educational applications, despite contrary pressures to use AI as an alternative to human decision making. Policies should not hinder innovation and improvement, nor should they be burdensome to implement. Society needs an education-focused AI policy that protects civil rights and promotes democratic values in the building, deployment, and governance of automated systems to be used across the many decentralized levels of the American educational system.

Foundation 2: Advance Equity

“AI brings educational technology to an inflection point. We can either increase disparities or shrink them, depending on what we do now.”
—Dr. Russell Shilling

A recent Executive Order⁹ issued by President Biden sought to strengthen the connection among racial equity, education and AI, stating that “members of underserved communities—many of whom have endured generations of discrimination and disinvestment—still confront significant barriers to realizing the full promise of our great Nation, and the Federal Government has a responsibility to remove these barriers” and that the Federal Government shall both “pursue educational equity so that our Nation’s schools put every student on a path to success” and also “root out bias in the design and use of new technologies, such as artificial intelligence.” A specific vision of equity, such as described in the Department’s recent report, *Advancing Digital Equity for All*¹⁰ is essential to policy discussion about AI in education. This report defines digital equity as

⁸ The White House (September 8, 2022). Readout of White House listening session on tech platform accountability. <https://www.whitehouse.gov/briefing-room/statements-releases/2022/09/08/readout-of-white-house-listening-session-on-tech-platform-accountability/>

⁹ The White House (February 17, 2023). Executive order on further advancing racial equity and support for underserved communities through the federal government. <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/02/16/executive-order-on-further-advancing-racial-equity>

¹⁰ U.S. Department of Education, Office of Educational Technology (2022). Advancing digital equity for all: Community-based recommendations for developing effective digital equity plans to close the digital divide and enable technology-empowered learning. US Department of Education.

“the condition in which individuals and communities have the information technology capacity that is needed for full participation in the society and economy of the United States.”


Issues related to racial equity and unfair bias were at the heart of every listening session we held. In particular, we heard a conversation that was increasingly attuned to issues of data quality and the consequences of using poor or inappropriate data in AI systems for education. Datasets are used to develop AI, and when they are non-representative or contain undesired associations or patterns, resulting AI models may act unfairly in how they detect patterns or automate decisions. Systematic, unwanted unfairness in how a computer detects patterns or automates decisions is called “algorithmic bias.” Algorithmic bias could diminish equity at scale with unintended discrimination. As this document discussed in the *Formative Assessment* section, this is not a new conversation. For decades, constituents have rightly probed whether assessments are unbiased and fair. Just as with assessments, whether an AI model exhibits algorithmic bias or is judged to be fair and trustworthy is critical as local school leaders make adoption decisions about using AI to achieve their equity goals.

We highlight the concept of “algorithmic discrimination” in the Blueprint. Bias is intrinsic to how AI algorithms are developed using historical data, and it can be difficult to anticipate all impacts of biased data and algorithms during system design. The Department holds that biases in AI algorithms must be addressed when they introduce or sustain unjust discriminatory practices in education. For example, in postsecondary education, algorithms that make enrollment decisions, identify students for early intervention, or flag possible student cheating on exams must be interrogated for evidence of unfair discriminatory bias—and not only when systems are designed, but also later, as systems become widely used.

Foundation 3: Ensure Safety, Ethics, and Effectiveness

A central safety argument in the Department’s policies is the need for data privacy and security in the systems used by teachers, students, and others in educational institutions. The development and deployment of AI requires access to detailed data. This data goes beyond conventional student records (roster and gradebook information) to detailed information about what students do as they learn with technology and what teachers do as they use technology to teach. AI’s dependence on data requires renewed and strengthened attention to data privacy, security, and governance (as also indicated in the Blueprint). As AI models are not generally developed in consideration of educational usage or student privacy, the educational application of these models may not be aligned with the educational institution’s efforts to comply with federal student privacy laws, such as FERPA, or state privacy laws.

Figure 2: The Elementary and Secondary Education Act defines four levels of evidence.



Further, educational leaders are committed to basing their decisions about the adoption of educational technology on evidence of effectiveness—a central foundation of the Department’s policy. For example, the requirement to base decisions on evidence also arises in the Elementary and Secondary Education Act (ESEA), as amended, which introduced four tiers of evidence (see Figure 2). Our nation’s research agencies, including the Institute of Education Sciences, are essential to producing the needed evidence. The Blueprint calls for evidence of effectiveness, but the education sector is ahead of that game: we need to insist that AI-enhanced edtech rises to meet ESEA standards as well.

Foundation 4: Promote Transparency

The central role of complex AI models in a technology’s detection of patterns and implementation of automation is an important way in which AI-enabled applications, products, and services will be different from conventional edtech. The Blueprint introduces the need for transparency about AI models in terms of disclosure (“notice”) and explanation. In education, decision makers will need more than notice—they will need to understand how AI models work in a range of general educational use cases, so they can better anticipate limitations, problems, and risks.

AI models in edtech will be approximations of reality and, thus, constituents can always ask these questions: **How precise are the AI models? Do they accurately capture what is most important? How well do the recommendations made by an AI model fit educational goals? What are the broader implications of using AI models at scale in educational processes?**

Building on what was heard from constituents, the sections of this report develop the theme of evaluating the quality of AI systems and tools using multiple dimensions as follows:

- **About AI:** AI systems and tools must respect data privacy and security. Humans must be in the loop.
- **Learning:** AI systems and tools must align to our collective vision for high-quality learning, including equity.
- **Teaching:** AI systems and tools must be inspectable, explainable, and provide human alternatives to AI-based suggestions; educators will need support to exercise professional judgment and override AI models, when necessary.

- **Formative Assessment:** AI systems and tools must minimize bias, promote fairness, and avoid additional testing time and burden for students and teachers.
- **Research and Development:** AI systems and tools must account for the context of teaching and learning and must work well in educational practice, given variability in students, teachers, and settings.
- **Recommendations:** Use of AI systems and tools must be safe and effective for students. They must include algorithmic discrimination protections, protect data privacy, provide notice and explanation, and provide a recourse to humans when problems arise. The people most affected by the use of AI in education must be part of the development of the AI model, system, or tool, even if this slows the pace of adoption.

We return to the idea that these considerations fit together in a comprehensive perspective on the quality of AI models in the *Recommendations* section.

Overview of Document

We begin in the next section by elaborating a definition of AI, followed by addressing learning, teaching, assessment, and research and development. Organizing key insights by these topics keeps us focused on exploring implications for improving educational opportunity and outcomes for students throughout the report.

Within these topics, three important themes are explored:

1. **Opportunities and Risks.** Policies should focus on the most valuable educational advances while mitigating risks.
2. **Trust and Trustworthiness.** Trust and safeguarding are particularly important in education because we have an obligation to keep students out of harm's way and safeguard their learning experiences.
3. **Quality of AI Models.** The process of developing and then applying a model is at the heart of any AI system. Policies need to support evaluation of the qualities of AI models and their alignment to goals for teaching and learning during the processes of educational adoption and use.


"AI in education can only grow at the speed of trust."
—Dr. Dale Allen

What is AI?

Our preliminary definition of AI as automation based on associations requires elaboration. Below we address three additional perspectives on what constitutes AI. Educators will find these different perspectives arise in the marketing of AI functionality and are important to understand when evaluating edtech systems that incorporate AI. One useful glossary of AI for Education terms is the [CIRCLS Glossary of Artificial Intelligence Terms for Educators](#).¹¹

AI is not one thing but an umbrella term for a growing set of modeling capabilities, as visualized in Figure 3.

Figure 3: Components, types, and subfields of AI based on Regona et al (2022).¹²



¹¹ Search for “AI Glossary Educators” to find other useful definitions.

¹² Regona, Massimo & Yigitcanlar, Tan & Xia, Bo & Li, R.Y.M. (2022). Opportunities and adoption challenges of AI in the construction industry: A PRISMA review. *Journal of Open Innovation Technology Market and Complexity*, 8(45). <https://doi.org/10.3390/joitmc8010045>

Perspective: Human-Like Reasoning

"The theory and development of computer systems able to perform tasks normally requiring human intelligence such as, visual perception, speech recognition, learning, decision-making, and natural language processing." ¹³

Broad cultural awareness of AI may be traced to the landmark 1968 film "2001: A Space Odyssey"—in which the "Heuristically-programmed ALgorithmic" computer, or "HAL," converses with astronaut Frank. HAL helps Frank pilot the journey through space, a job that Frank could not do on his own. However, Frank eventually goes outside the spacecraft, HAL takes over control, and this does not end well for Frank. HAL exhibits human-like behaviors, such as reasoning, talking, and acting. Like all applications of AI, HAL can help humans but also introduces unanticipated risks—especially since AI reasons in different ways and with different limitations than people do.

The idea of "human-like" is helpful because it can be a shorthand for the idea that computers now have capabilities that are very different from the capabilities of early edtech applications. Educational applications will be able to converse with students and teachers, co-pilot how activities unfold in classrooms, and take actions that impact students and teachers more broadly. There will be both opportunities to do things much better than we do today and risks that must be anticipated and addressed.

The "human-like" shorthand is not always useful, however, because AI processes information differently from how people process information. When we gloss over the differences between people and computers, we may frame policies for AI in education that miss the mark.

Perspective: An Algorithm that Pursues a Goal

"Any computational method that is made to act independently towards a goal based on inferences from theory or patterns in data." ¹⁴

This second definition emphasizes that AI systems and tools identify patterns and choose actions to achieve a given goal. These pattern recognition capabilities and automated recommendations will be used in ways that impact the educational process, including student learning and teacher instructional decision making. For example, today's personalized learning systems may recognize signs that a student is struggling and may recommend an alternative instructional sequence. The scope of pattern recognition and automated recommendations will expand.

¹³ IEEE-USA Board of Directors. (February 10, 2017). *Artificial intelligence research, development and regulation*. IEEE <http://globalpolicy.ieee.org/wp-content/uploads/2017/10/IEEE17003.pdf>

¹⁴ Friedman, L., Blair Black, N., Walker, E., & Roschelle, J. (November 8, 2021) *Safe AI in education needs you*. Association of Computing Machinery blog, <https://cacm.acm.org/blogs/blog-cacm/256657-safe-ai-in-education-needs-you/fulltext>

Correspondingly, humans must determine the types and degree of responsibility we will grant to technology within educational processes, which is not a new dilemma.

For decades, the lines between the role of teachers and computers have been discussed in education, for example, in debates using terms such as “computer-aided instruction,” “blended instruction,” and “personalized learning.” Yet, how are instructional choices made in systems that include both humans and algorithms? Today, AI systems and tools are already enabling the adaptation of instructional sequences to student needs to give students feedback and hints, for example, during mathematics problem solving or foreign language learning. This discussion about the use of AI in classroom pedagogy and student learning will be renewed and intensify as AI-enabled systems and tools advance in capability and become more ubiquitous.

Let’s start with another simple example. When a teacher says, “Display a map of ancient Greece on the classroom screen,” an AI system may choose among hundreds of maps by noting the lesson objectives, what has worked well in similar classrooms, or which maps have desirable features for student learning. In this case, when an AI system suggests an instructional resource or provides a choice among a few options, the instructor may save time and may focus on more important goals. However, there are also forms of AI-enabled automation that the classroom instructor may reject, for example, enabling an AI system or tool to select the most appropriate and relevant readings for students associated with a historical event. In this case, an educator may choose not to utilize AI-enabled systems or tools given the risk of AI creating false facts (“hallucinating”) or steering students toward inaccurate depictions of historical events found on the internet. Educators will be weighing benefits and risks like these daily.

Computers process theory and data differently than humans. AI’s success depends on associations or relationships found in the data provided to an algorithm during the AI model development process. Although some associations may be useful, others may be biased or inappropriate. Finding bad associations in data is a major risk, possibly leading to algorithmic discrimination. Every guardian is familiar with the problem: A person or computer may say, “Our data suggests your student should be placed in this class,” and the guardian may well argue, “No, you are using the wrong data. I know my child better, and they should instead be placed in another class.” This problem is not limited exclusively to AI systems and tools, but the use of AI models can amplify the problem when a computer uses data to make a recommendation because it may appear to be more objective and authoritative, even if it is not.

Although this perspective can be useful, it can be misleading. A human view of agency, pursuing goals, and reasoning includes our human abilities to make sense of multiple contexts. For example, a teacher may see three students each make the same mathematical error but recognize that one student has an Individualized Education Program to address vision issues, another misunderstands a mathematical concept, and a third just experienced a frustrating interaction on the playground; the same instructional decision is therefore not appropriate. However, AI systems often lack data and judgement to appropriately include context as they detect patterns and automate decisions. Further, case studies show that technology has the potential to quickly derail from safe to unsafe or from effective to ineffective when the context shifts even slightly. For this and other reasons, people must be involved in goal setting, pattern analysis, and decision-making.¹⁵

¹⁵ Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking. ISBN 978-0-525-55861-3.

Perspective: Intelligence Augmentation

"Augmented intelligence is a design pattern for a human-centered partnership model of people and artificial intelligence (AI) working together to enhance cognitive performance, including learning, decision making, and new experiences." ¹⁶

Foundation #1 (above) keeps humans in the loop and positions AI systems and tools to support human reasoning. "Intelligence Augmentation" (IA)¹⁷ centers "intelligence" and "decision making" in humans but recognizes that people sometimes are overburdened and benefit from assistive tools. AI may help teachers make better decisions because computers notice patterns that teachers can miss. For example, when a teacher and student agree that the student needs reminders, an AI system may provide reminders in whatever form a student likes without adding to the teacher's workload. Intelligence Automation (IA) uses the same basic capabilities of AI, employing associations in data to notice patterns, and, through automation, takes actions based on those patterns. However, IA squarely focuses on helping people in human activities of teaching and learning, whereas AI tends to focus attention on what computers can do.

Definition of "Model"

The above perspectives open a door to making sense of AI. Yet, to assess AI meaningfully, constituents must consider specific models and how they are developed. In everyday usage, the term "model" has multiple definitions. We clarify our intended meaning, which is a meaning similar to "mathematical model," below. (Conversely, note that "model" as used in "AI model" is unlike the usage in "model school" or "instructional model" as AI model is not a singular case created by experts to serve as an exemplar.)

AI models are like financial models: an approximation of reality that is useful for identifying patterns, making predictions, or analyzing alternative decisions. In a typical middle school math curriculum, students use a mathematical model to analyze which of two cell phone plans is better. Financial planners use this type of model to provide guidance on a retirement portfolio. At its heart, AI is a highly advanced mathematical toolkit for building and using models. Indeed, in well-known chatbots, complex essays are written one word at a time. The underlying AI model predicts which next words would likely follow the text written so far; AI chatbots use a very large statistical model to add one likely word at a time, thereby writing surprisingly coherent essays.

When we ask about the model at the heart of AI, we begin to get answers about "what aspects of reality does the model approximate well?" and "how appropriate is it to the decision to be made?" One could similarly ask about algorithms—the specific decision-making processes that an AI model uses to go from inputs to outputs. One could also ask about the quality of the data used to build the model—for example, how representative is that data? Switching among three terms—

¹⁶ Gartner (n.d.) *Gartner glossary: Augmented intelligence*. Gartner. <https://www.gartner.com/en/information-technology/glossary/augmented-intelligence>

¹⁷ Englebart, D.C. (October 1962). Augmenting human intellect: A conceptual framework. SRI Summary Report AFOSR-3223. <https://www.dougenelbart.org/pubs/augment-3906.html>

models, algorithms, and data—will become confusing. Because the terms are closely related, we've chosen to focus on the concept of AI models. We want to bring to the fore the idea that every AI model is incomplete, and it's important to know how well the AI model fits the reality we care about, where the model will break down, and how.

Sometimes people avoid talking about the specifics of models to create a mystique. Talking as though AI is unbounded in its potential capabilities and a nearly perfect approximation to reality can convey an excitement about the possibilities of the future. The future, however, can be oversold. Similarly, sometimes people stop calling a model AI when its use becomes commonplace, yet such systems are still AI models with all of the risks discussed here. We need to know exactly when and where AI models fail to align to visions for teaching and learning.

Insight: AI Systems Enable New Forms of Interaction

AI models allow computational processes to make recommendations or plans and also enable them to support forms of interaction that are more natural, such as speaking to an assistant. AI-enabled educational systems will be desirable in part due to their ability to support more natural interactions during teaching and learning. In classic edtech platforms, the ways in which teachers and students interact with edtech are limited. Teachers and students may choose items from a menu or in a multiple-choice question. They may type short answers. They may drag objects on the screen or use touch gestures. The computer provides outputs to students and teachers through text, graphics, and multimedia. Although these forms of inputs and outputs are versatile, no one would mistake this style of interaction with the way two people interact with one another; it is specific to human-computer interaction. With AI, interactions with computers are likely to become more like human-to-human interactions (see Figure 4). A teacher may speak to an AI assistant, and it may speak back. A student may make a drawing, and the computer may highlight a portion of the drawing. A teacher or student may start to write something, and the computer may finish their sentence—as when today's email programs can complete thoughts faster than we can type them.

Additionally, the possibilities for automated actions that can be executed by AI tools are expanding. Current personalization tools may automatically adjust the sequence, pace, hints, or trajectory through learning experiences.¹⁸ Actions in the future might look like an AI system or tool that helps a student with homework¹⁹ or a teaching assistant that reduces a teacher's workload by recommending lesson plans that fit a teacher's needs and are similar to lesson plans a teacher previously liked.²⁰ Further, an AI-enabled assistant may appear as an additional “partner” in a small group of students who are working together on a collaborative assignment.²¹ An AI-enabled tool may also help teachers with complex classroom routines.²² For example, a

¹⁸ Shemshack, A., Spector, J.M. (2020) A systematic literature review of personalized learning terms. *Smart Learning Environments*, 7(33). <https://doi.org/10.1186/s40561-020-00140-9>

¹⁹ Roschelle, J., Feng, M., Murphy, R. & Mason, C.A. (2016). Online mathematics homework increases student achievement. *AERA Open*, 2(4), 1-12. DOI: [10.1177/2332858416673968](https://doi.org/10.1177/2332858416673968)

²⁰ Celik, I., Dindar, M., Muukkonen, H. & Järvelä, S. (2022). The promises and challenges of artificial intelligence for teachers: A systematic review of research. *TechTrends*, 66, 616–630. <https://doi.org/10.1007/s11528-022-00715-y>

²¹ Chen, C., Park, H.W. & Breazeal, C. (2020). Teaching and learning with children: Impact of reciprocal peer learning with a social robot on children's learning and emotive engagement. *Computers & Education*, 150, <https://doi.org/10.1016/j.compedu.2020.103836>

²² Holstein, K., McLaren, B.M., & Aleven, V. (2019). Co-designing a real-time classroom orchestration tool to support teacher-AI complementarity. *Journal of Learning Analytics*, 6(2). <https://doi.org/10.18608/jla.2019.62.3>

tool may help teachers with orchestrating²³ the movement of students from a full class discussion into small groups and making sure each group has the materials needed to start their work.

Figure 4. Differences that teachers and students may experience in future technologies.

	Familiar Technology Capabilities	Future Technology Capabilities
Input	• Typing	• Speaking
	• Clicking and dragging	• Drawing
	• Touching and gesturing	• Analyzing images and video
Processing	• Displaying information and tasks	• Assisting students and teachers
	• Sequencing learning activities	• Planning and adapting activities
	• Checking student work	• Revealing patterns in student work
Output	• Text	• Conversations
	• Graphics	• Annotating and highlighting
	• Multimedia	• Suggesting and recommending
	• Dashboards	• Organizing and guiding

Key Recommendation: Human in the Loop AI

Many have experienced a moment where technology surprised them with an uncanny ability to recommend what feels like a precisely personalized product, song, or even phrase to complete a sentence in a word processor such as the one being used to draft this document. Throughout this supplement, we talk about specific, focused applications where AI systems may bring value (or risks) into education. At no point do we intend to imply that AI can replace a teacher, a guardian, or an educational leader as the custodian of their students' learning. We talk about the limitations of models in AI and the conversations that educational constituents need to have about what qualities they want AI models to have and how they should be used.

"We can use AI to study the diversity, the multiplicity of effective learning approaches and think about the various models to help us get a broader understanding of what effective, meaningful engagement might look like across a variety of different contexts."

—Dr. Marcelo Aaron Bonilla Worsley

²³ Roschelle, J., Dimitriadis, Y. & Hoppe, U. (2013). Classroom orchestration: Synthesis. *Computers & Education*, 69, 512-526. <https://doi.org/10.1016/j.compedu.2013.04.010>

These limitations lead to our first recommendation: that we pursue a vision of AI where humans are in the loop. That means that people are part of the process of noticing patterns in an educational system and assigning meaning to those patterns. It also means that teachers remain at the helm of major instructional decisions. It means that formative assessments involve teacher input and decision making, too. One loop is the cycle of recognizing patterns in what students do and selecting next steps or resources that could support their learning. Other loops involve teachers planning and reflecting on lessons. Response to Intervention is another well-known type of loop.

The idea of humans in the loop is part of our broader discussions happening about AI and society, not just AI in education. Interested readers could look for more on human-centered AI, responsible AI, value-sensitive AI, AI for social good, and other similar terms that ally with humans in the loop, such as “human-centered AI.”

Exercising judgement and control in the use of AI systems and tools is an essential part of providing the best opportunity to learn for all students—especially when educational decisions carry consequence. AI does not have the broad qualities of contextual judgment that people do. Therefore, people must remain responsible for the health and safety of our children, for all students’ educational success and preparation for their futures, and for creating a more equitable and just society.

Learning

The Department's long-standing edtech vision sees students as active learners; students participate in discussions that advance their understanding, use visualizations and simulations to explain concepts as they relate to the real world, and leverage helpful scaffolding and timely feedback as they learn. Constituents want technology to align to and build on these and other research-based understandings of how people learn. Educators can draw upon two books titled *How People Learn* and *How People Learn II* by the National Academies of Sciences, Engineering, and Medicine for a broad synthesis of what we know about learning.²⁴ As we shape AI-enhanced edtech around research-based principles, a key goal must be to strengthen and support learning for those who have experienced unfavorable circumstances for learning, such as caused by the COVID-19 pandemic or by broader inequities. And we must keep a firm eye toward the forms of learning that will most benefit learners in their future lives in communities and workplaces.

Examples of AI supporting learning principles in this section include the following: AI-based tutoring for students as they solve math problems (based on cognitive learning theories), adapting to learners with special needs (based on the Universal Design for Learning framework and related theories), and AI support for effective student teamwork (based on theories in the field called "Computer Supported Collaborative Learning").

Insight: AI Enables Adaptivity in Learning

Adaptivity has been recognized as a key way in which technology can improve learning.²⁵ AI can be a toolset for improving the adaptivity of edtech. AI may improve a technology's ability to meet students where they are, build on their strengths, and grow their knowledge and skills. Because of AI's powers of work with natural forms of input and the foundational strengths of AI models (as discussed in the *What is AI?* section), AI can be an especially strong toolkit for expanding the adaptivity provided to students.

And yet, especially with AI, adaptivity is always more specific and limited than what a broad phrase like "meet students where they are" might suggest. Core limits arise from the nature of the model at the heart of any specific AI-enabled system. Models are approximations of reality. When important parts of human learning are left out of the model or less fully developed, the resulting adaptivity will also be limited, and the resulting supports for learning may be brittle or narrow. Consequently, this section on *Learning* focuses on one key concept: Work toward AI models that fit the fullness of visions for learning—and avoid limiting learning to what AI can currently model well.

AI models are demonstrating greater skills because of advances in what are called "large language models" or sometimes "foundational models." These very general models still have limits. For example, generative AI models discussed in the mainstream news can quickly generate convincing essays about a wide variety of topics while other models can draw credible images based on just a few prompts. Despite the excitement about foundational models, experts in our

²⁴ National Research Council. 2000. *How people learn: Brain, mind, experience, and school*. The National Academies Press. <https://doi.org/10.17226/9853>; National Academies of Sciences, Engineering, and Medicine. 2018. *How people learn II: Learners, contexts, and cultures*. The National Academies Press. <https://doi.org/10.17226/24783>

²⁵ Alevan, V., McLaughlin, E. A., Glenn, R. A., & Koedinger, K. R. (2016). Instruction based on adaptive learning technologies. In Mayer, R.E. & Alexander, P.A., *Handbook of research on learning and instruction*, 522-560. ISBN: 113883176X

listening sessions warned that AI models are narrower than visions for human learning and that designing learning environments with these limits in mind remains very important. The models are also brittle and can't perform well when contexts change. In addition, they don't have the same "common sense" judgment that people have, often responding in ways that are unnatural or incorrect.²⁶ Given the unexpected ways in which foundational models miss the mark, keeping humans in the loop remains highly important.

Intelligent Tutoring Systems: An Example of AI Models

One long-standing type of AI-enabled technology is an Intelligent Tutoring System (ITS).²⁷ In an early success, scientists were able to build accurate models of how human experts solve mathematical problems. The resulting model was incorporated into a system that would observe student problem solving as they worked on mathematical problems on a computer. Researchers who studied human tutors found that feedback on specific steps (and not just right or wrong solutions) is a likely key to why tutoring is so effective.²⁸ For example, when a student diverged from the expert model, the system gave feedback to help the student get back on track.²⁹ Importantly, this feedback went beyond right or wrong, and instead, the model was able to provide feedback on specific steps of a solution process. A significant advancement of AI, therefore, can be its ability to provide adaptivity at the step-by-step level and its ability to do so at scale with modest cost.

As a research and development (R&D) field emerged to advance ITS, the work has gone beyond mathematics problems to additional important issues beyond step-by-step problem solving. In the early work, some limitations can be observed. The kinds of problems that an ITS could support were logical or mathematical, and they were closed tasks, with clear expectations for what a solution and solution process should look like. Also, the "approximation of reality" in early AI models related to cognition and not to other elements of human learning, for example, social or motivational aspects. Over time, these early limitations have been addressed in two ways: by expanding the AI models and by involving humans in the loop, a perspective that is also important now. Today, for example, if an ITS specializes in feedback as a student practices, a human teacher could still be responsible for motivating student engagement and self-regulation along with other aspects of instruction. In other contemporary examples, the computer ITS might focus on problem solving practice, while teachers work with students in small groups. Further, students can be in the loop with AI, as is the case with "open learner models"—a type of AI-enabled system that provides information to support student self-monitoring and reflection.³⁰

²⁶ Dieterle, E., Dede, C. & Walker, M. (2022). The cyclical ethical effects of using artificial intelligence in education. *AI & Society*. <https://link.springer.com/article/10.1007/s00146-022-01497-w>

²⁷ Mousavinasab, E., Zarifsanaiy, N., R. Niakan Kalhori, S., Rakhshan, M., Keikha, L., & Ghazi Saeedi, M. (2021). Intelligent tutoring systems: A systematic review of characteristics, applications, and evaluation methods. *Interactive Learning Environments*, 29(1), 142–163. <https://psycnet.apa.org/doi/10.1080/10494820.2018.1558257>

²⁸ Van Lehn, K. (2011) The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221. <https://doi.org/10.1080/00461520.2011.611369>

²⁹ Ritter, S., Anderson, J.R., Koedinger, K.R. & Corbett, A. (2007). Cognitive Tutor: Applied research in mathematics education. *Psychonomic Bulletin & Review*, 14, 249–255/ <https://doi.org/10.3758/BF03194060>

³⁰ Winne, P.H. (2021). Open learner models working in symbiosis with self-regulating learners: A research agenda. *International Journal of Artificial Intelligence in Education*, 31(3), 446–459. <https://doi.org/10.1007/s40593-020-00212-4>

Although R&D along the lines of an ITS should not limit the view of what's possible, such an example is useful because so much research and evaluation has been done on the ITS approach. Researchers have looked across all the available high-quality studies in a meta-analysis and concluded that ITS approaches are effective.³¹ Right now, many school systems are looking at high-intensity human tutoring to help students with unfinished learning. Human tutoring is very expensive, and it is hard to find enough high-quality human tutors. With regard to large-scale needs, if it is possible for an ITS to supplement what human tutors do, it might be possible to extend beyond the amount of tutoring that people can provide to students.

Important Directions for Expanding AI-Based Adaptivity

Adaptivity is sometimes referred to as “personalization.” Although this is a convenient term, many observers have noted how imprecise it is.³² For some educators, personalization means giving learners “voice and choice,” and for others it means that a learning management system recommends an individual “playlist” of activities to each student. Hidden in that imprecision is the reality that many edtech products that personalize do so in limited ways. Adjusting the difficulty and the order of lesson materials are among the two most common ways that edtech products adapt. And yet, any teacher knows there is more to supporting learning than adjusting the difficulty and sequence of materials. For example, a good teacher can find ways to engage a student by connecting to their own past experiences and can shape explanations until they really connect in an “aha!” moment for that student. When we say, “meet the learner where they are,” human teachers bring a much more complete picture of each learner than most available edtech. The teacher is also not likely to “over personalize” (by performing like an algorithm that only presents material for which the learner has expressed interest), thereby limiting the student’s exposure to new topics. The nature of “teachable moments” that a human teacher can grasp is broader than the teachable moments today’s AI models grasp.

In our listening sessions, we heard many ways in which the core models in an AI system must be expanded. We discuss these below.

1. **From deficit-based to asset-oriented.** Listening session attendees noted that the rhetoric around adaptivity has often been deficit-based; technology tries to pinpoint what a student is lacking and then provides instruction to fill that specific gap. Teachers also orient to students’ strengths; they find competencies or “assets” a student has and use those to build up the students’ knowledge. AI models cannot be fully equitable while failing to recognize or build upon each student’s sources of competency. AI models that are more asset-oriented would be an advance.
2. **From individual cognition to including social and other aspects of learning.** The existing adaptivity rhetoric has also tended to focus on individualized learning and mostly on cognitive elements of learning, with motivational and other elements only brought in to support the cognitive learning goals. Attendees observe that their vision for learning is broader than cognition. Social learning is important, for example, especially

³¹ Kulik, J.A., & Fletcher, J.D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*, 86(1), 42–78; Ma, W., Adesope, O.O., Nesbit, J.C. & Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology*, 106(4), 901–918. <http://dx.doi.org/10.1037/a0037123>

³² Plass, J.L., & Pawar, S. (2020). Toward a taxonomy of adaptivity for learning. *Journal of Research on Technology in Education*, 52(3), 275–300. <https://doi.org/10.1080/15391523.2020.1719943>

for students to learn to reason, explain, and justify. For students who are learning English, customized and adaptive support for improving language skills while learning curricular content is clearly important. Developing self-regulation skills is also important. A modern vision of learning is not individualistic; it recognizes that students learn in groups and communities too.

3. **From neurotypical to neurodiverse learners.** AI models could help in including neurodiverse learners (students who access, process, and interact with the world in less common ways than “neurotypical” students) who could benefit from different learning paths and from forms of display and input that fit their strengths. Constituents want AI models that can support learning for neurodiverse learners and learners with disabilities. Thus, they want AI models that can work with multiple paths to learning and multiple modalities of interaction. Such models should be tested for efficacy, to guard against the possibility that some students could be assigned a “personalized” but inadequate learning resource. In addition, some systems for neurodiverse students are presently underutilized, so designs that support intended use will also be important.
4. **From fixed tasks to active, open, and creative tasks.** As mentioned above, AI models are historically better at closed tasks like solving a math problem or logical tasks like playing a game. In terms of life-wide and lifelong opportunities, we value learning how to succeed at open-ended and creative tasks that require extended engagement from the learner, and these are often not purely mathematical or logical. We want students to learn to invent and create innovative approaches. We want AI models that enable progress on open, creative tasks.
5. **From correct answers to additional goals.** At the heart of many adaptivity approaches now on the market, the model inside the technology counts students' wrong answers and decides whether to speed up, slow down, or offer a different type of learning support. Yet, right and wrong answers are not the only learning goals. We want students to learn how to self-regulate when they experience difficulties in learning, for example, such as being able to persist in working on a difficult problem or knowing how and when to ask for help. We want learners to become skilled in teamwork and in leading teams. As students grow, we want them to develop more agency and to be able to act on their own to advance toward their own learning goals.

Listing every dimension of expansion that we heard in our listening sessions is beyond the scope of this report. Some additional dimensions are presented in the following sections on *Teaching*, *Assessment*, and *Research*. For example, in *Research*, we discuss all the ways in which AI systems have trouble with context—context that humans readily grasp and consider.

Overall, constituents in the listening sessions realized we need an ambitious outlook on learning to respond to the future today's learners face. Constituents were concerned about ways in which AI might narrow learning. For example, if the incorporation of AI into education slowed attention to students' skills on creative, open-ended tasks and their ability to lead and collaborate in teams, then school districts may be less able to realize their students' progress in relation to a Portrait of a Graduate who excels in communication and other skills valued in communities and careers.

Constituents reminded us that as we conceptualize what we want AI in edtech to accomplish, we must start and constantly revisit a human-centered vision of learning.

A Duality: Learning With and About AI

As AI is brought into schools, two broad perspectives about AI in education arise: (1) AI in support of student learning; and (2) support for learning about AI and related technologies. So far, we've discussed AI systems and tools to support student learning and mastery of subjects like mathematics and writing. Yet, it is also important that students learn about AI, critically examine its presence in education and society, and determine its role and value in their own lives and careers. We discuss risks across each section in this report. Here, it is important for students to become more aware of and savvy to the risks of AI—including risks of bias and surveillance—as they appear in all elements of their lives. In the recent past, schools have supported students' understanding of cybersecurity, for example. AI will bring new risks, and students need to learn about them.

We are encouraged by efforts we've seen underway that would give students opportunities to learn about how AI works while also giving them opportunities to discuss relevant topics like privacy and security.³³ Other learning goals are noted in the [K-12 Computer Science Framework](#). We've seen that students can begin learning about AI in elementary, middle, and high school. They can use AI to design simulations and products that they find exciting. And we've seen that students want to talk about the ethics of products they experience in their everyday lives and have much to say about the kinds of products they'd like to see or not see in school. (And later, in the *Research* section, we note the desire for co-design processes that involve students in creating the next generation of AI-enabled edtech). Overall, it's important to balance attention to using AI to support learning and giving students opportunities to learn about AI.

A Challenge: Systems Thinking About AI in Education

As AI expands into the educational system, our listening session attendees reminded us that it will be entering parts or locations of the system that are presently dysfunctional. AI is certainly not a fix for broken systems, and instead, must be used with even more care when the systems' context is unstable or uncertain.

³³ Forsyth, S., Dalton, B., Foster, E.H., Walsh, B., Smilack, J., & Yeh, T. (2021, May). Imagine a more ethical AI: Using stories to develop teens' awareness and understanding of artificial intelligence and its societal impacts. In 2021 Conference on Research in Equitable and Sustained Participation in Engineering, Computing, and Technology (RESPECT). IEEE. <https://doi.org/10.1109/RESPECT51740.2021.9620549>; Zhang, H., Lee, I., Ali, S., DiPaola, D., Cheng, Y., & Breazeal, C. (2022). Integrating ethics and career futures with technical learning to promote AI literacy for middle school students: An exploratory study. *International Journal of Artificial Intelligence in Education*, 1–35. <https://doi.org/10.1007/s40593-022-00293-3>

"First and foremost, they are getting deployed in educational contexts that are already fragmented and broken and unequal. Technology doesn't discriminate—we do. So, as we think about the application of these new systems, we have to really think about the contextual application of AI."

—Dr. Nicole Turner

As discussed previously, because AI systems and tools do not fully align with goals for learning, we have to design educational settings to situate AI in the right place, where educators and other adults can make effective use of these tools for teaching and learning. Within the ITS example, we saw that AI could make learning by practicing math problems more effective, and a whole curricular approach might include roles for teachers that emphasize mathematical practices like argumentation and modeling. Further, small-group work is likely to remain important: Students might work in small groups to use mathematics to predict or justify as they work on responding to a realistic challenge. At the present, one “right place” for people, and not AI, is understanding how learning can be culturally responsive and culturally sustaining, as AI is not even close to being ready to connect learning to the unique strengths in a student’s community and family.

Open Questions About AI for Learning

With advances occurring in the foundations for AI, opportunities to use AI in support of learning are rapidly expanding. As we explore these opportunities, the open questions below deserve ongoing attention:

- To what extent is AI enabling adaptation to students’ strengths and not just deficits? Is AI enabling improved support for learners with disabilities and English language learners?
- How are youth voices involved in choosing and using AI for learning?
- Is AI leading to narrower student activities (e.g., procedural math problems), or the fuller range of activities highlighted in the National Educational Technology Plan (NETP), which emphasizes features such as personalized learning, project-based learning, learning from visualizations, simulations, and virtual reality, as well as learning across school, community, and familial settings?
- Is AI supporting the whole learner, including social dimensions of learning such as enabling students to be active participants in small group and collaborative learning? For example, does AI contribute to aspects of student collaboration we value like shared attention, mutual engagement, peer help, self-regulation, and building on each other’s contributions?
- When AI is used, are students’ privacy and data protected? Are students and their guardians informed about what happens with their data?
- How strong are the processes or systems for monitoring student use of AI for barriers, bias, or other undesirable consequences of AI use by learners? How are emergent issues addressed?
- Is high-quality research or evaluations about the impacts of using the AI system for student learning available? Do we know not only whether the system works but for whom and under what conditions?

Key Recommendation: Seek AI Models Aligned to a Vision for Learning

We've called attention to how advances in AI are important to adaptivity but also to ways in which adaptivity is limited by the model's inherent quality. We noted that a prior wave of edtech used the term "personalized" in differing ways, and it was often important to clarify what personalization meant for a particular product or service. Thus, our key recommendation is to tease out the strengths and limitations of AI models inside forthcoming edtech products and to focus on AI models that align closely to desired visions of learning. AI is now advancing rapidly, and we should differentiate between products that have simple AI-like features inside and products that have more sophisticated AI models.

Looking at what's happening in research and development, we can see significant effort and push toward overcoming these limitations. We noted that decision makers need to be careful about selecting AI models that might narrow their vision for learning, as general artificial intelligence does not exist. And because AI models will always be narrower than real world experience, we need to proceed with systems thinking in which humans are in the loop, with the strengths and weaknesses of the specific educational system considered. We hold that the full system for learning is broader than its AI component.

Teaching

Teachers have long envisioned many things that technology could make possible for teachers, their classrooms, and their students but not the changes wrought by the recent pandemic. Today, nearly all teachers have experienced uses of technologies for instruction that no one anticipated. Some of those experiences were positive, and others were not. All of the experiences provide an important context as we think further about teaching and technology.

There is a critical need to focus on addressing the challenges teachers experience. It must become easier for teachers to do the amazing work they always do. We must also remember why people choose the teaching profession and ensure they can do the work that matters. This section discusses examples of AI supporting teachers and teaching including these concepts: AI assistants to reduce routine teaching burdens; AI that provides teachers with recommendations for their students' needs and extends their work with students; and AI that helps teachers to reflect, plan, and improve their practice.

"One opportunity I see with AI is being able to reduce the amount of attention I have to give to administrative things and increase the amount of attention I can give to my students with their learning needs in the classroom. So that's the first one that I'd say that I'm super excited about the possibility of AI to support me as a teacher."

—Vidula Plante

Always Center Educators in Instructional Loops


To succeed with AI as an enhancement to learning and teaching, we need to always center educators (ACE). Practically speaking, practicing "ACE in AI" means keeping a humanistic view of teaching front and center. ACE leads the Department to confidently respond "no" when asked "will AI replace teachers?" ACE is not just about making teachers' jobs easier but also making it possible to do what most teachers want to do. That includes, for example, understanding their students more deeply and having more time to respond in creative ways to teachable moments.

To bring more precision to how and where we should center educators, we return to our advocacy for human in the loop AI and ask, what are the loops in which teachers should be centered? Figure 5 suggests three key loops (inspired by research on adaptivity loops³⁴):

³⁴ Aleven, V., McLaughlin, E.A., Glenn, R.A., & Koedinger, K.R. (2016). Instruction based on adaptive learning technologies. In Mayer, R.E. & Alexander, P.A., *Handbook of research on learning and instruction*, 522-560. ISBN: 113883176X

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

Figure 5: Three ways to center educators as we conceptualize human in the loop AI



Please note that in the next section, on *Formative Assessment*, we also discuss teachers' important role in feedback loops that support students and enable school improvement. That section also includes a discussion of the concepts of "bias" and "fairness," which are important to teachers.

Insight: Using AI to Improve Teaching Jobs


The job of teaching is notoriously complex, with teachers making thousands of decisions each day. Teachers participate in classroom processes, in interactions with students beyond classrooms, in work with fellow teachers, and in administrative functions. They also are part of their communities and thus are expected to interact with families and caregivers.

If the teacher is able to efficiently predict and understand the range of other answers given by students in the class, it becomes possible to think creatively about the novel answer and figure how and why the student might have generated it.³⁵

We think about how much easier some everyday tasks have become. We can request and receive alerts and notifications about events. Selecting music that we want to hear used to be a multistep process (even with digital music), and now we can speak the name of a song we want to hear, and it plays. Likewise, mapping a journey used to require a cumbersome study of maps, but now cell phones let us choose among several transportation options to reach a destination. Why can't teachers be supported to notice changing student needs and provided with supports to enact a technology-rich lesson plan? Why can't they more easily plan their students' learning journeys? When things change in a classroom, as they always do, why don't the tools of the classroom make it easier for teachers to adapt to student strengths and needs on the fly?

Figure 6: Teachers work about 50 hours a week, spending less than half the time in direct interaction with students.

Activity composition of teacher working hours, number of hours



¹ Average for respondents in Canada, Singapore, United Kingdom, and United States.

² Includes a small "other" category.

Source: McKinsey Global Teacher and Student Survey

A report by McKinsey³⁶ first suggested that AI's initial benefit could be to improve teaching jobs by reducing low-level burdens in administrative or clerical work (Figure 6). The report also suggests that recovered time from AI-enabled technology should be rededicated toward more

³⁵ Hammerness, K., Darling-Hammond, L., & Bransford, J. (2005). *Preparing teachers for a changing world: What teachers should learn and be able to do*. Jossey-Bass. ISBN: 0787996343

³⁶ Bryant, J., Heitz, C., Sanghvi, S., & Wagle, D. (2020, January 14). *How artificial intelligence will impact K-12 teachers*. McKinsey. <https://www.mckinsey.com/industries/education/our-insights/how-artificial-intelligence-will-impact-k-12-teachers>

effective instruction—particularly, outcomes such as reducing the average 11 hours of weekly preparation down to only six. We highlight these opportunities and two others below.

1. **Handling low-level details to ease teaching burdens and increase focus on students.** A good teacher must master all levels of details, big and small. When working with a particular student, the teacher may wish to later send that student a helpful learning resource. How will they remember to send it? A voice assistant or other forms of an AI assistant could make it easier to stay organized by categorizing simple voice notes for teachers to follow up on after a classroom session ends. We are beginning to see AI-enabled voice assistants in the market, and they could do many simple tasks so that the teachers can stay focused on students. These tasks can include record-keeping, starting and stopping activities, controlling displays, speakers, and other technologies in the classroom, and providing reminders. Many workers may eventually use assistants to make their jobs easier, and teachers are the most deserving of efforts to ease their jobs now.
2. **Extending beyond the teacher's availability with their students but continuing to deliver on the teacher's intent.** Teachers almost always want to do more with each student than they can, given the limited number of hours before the next school day. A teacher may wish to sit with the student as they practice 10 more math problems, giving them ongoing support and feedback. If the teacher can sit with the student for only three problems, perhaps they could delegate to an AI-enabled learning system to help with the rest. Teachers cannot be at their best if on call at all hours to help with homework, but perhaps they can indicate what types of supports, hints, and feedback they want students to receive while studying after school hours. An AI assistant can ensure that students have that support wherever and whenever they do homework or practice skills on their own. Teachers may wish to provide more extensive personal notes to families/caregivers, and perhaps an AI assistant could help with drafts based on students' recent classroom work. Then, the teacher could review the AI-generated comments and quickly edit where needed before returning it to the student for another draft. AI tools might also help teachers with language translation so they can work with all parents and caregivers of their students. AI tools might also help teachers with awareness. For example, in the next section, *Formative Assessment*, we note that teachers can't always know what's going on for each student and in each small group of students; emerging products might signal to the teacher when a student or teacher may need some more personal attention.
3. **Making teacher professional development more productive and fruitful.** Emerging products already enable a teacher to record her classroom and allow an AI algorithm to suggest highlights of the classroom discussion worth reviewing with a professional development coach.³⁷ AI can compute metrics, such as whether students have been talking more or less, which are difficult for a teacher to calculate during a lesson.³⁸ For


³⁷ Chen, G., Clarke, S., & Resnick, L.B. (2015). Classroom Discourse Analyzer (CDA): A discourse analytic tool for teachers. *Technology, Instruction, Cognition and Learning*, 10(2), 85-105

³⁸ Jensen, E., Dale, M., Donnelly, P.J., Stone, C., Kelly, S., Godley, A. & D'Mello, S.K. (2020). Toward automated feedback on teacher discourse to enhance teacher learning. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20). <https://doi.org/10.1145/3313831.3376418>

teachers who want to increase student engagement, these metrics can be a valuable tool. Classroom simulation tools are also emerging and can enable teachers to practice their skills in realistic situations.³⁹ Simulators can include examples of teaching from a real classroom while changing the faces and voices of the participants so that teaching situations can be shared and discussed among teachers without revealing identities.

Note the emphasis above on what listening-session panelist Sarah Hampton said about the human touch. Teachers will feel that AI is helping them teach with a focus on their human connection to their students when the necessary (but less meaningful) burdens of teaching are lessened. In Figure 7, below, see concerns that teachers raised about AI during listening sessions.

Figure 7: Concerns raised during the listening session about teaching with AI



Preparing and Supporting Teachers in Planning and Reflecting

ACE also means preparing teachers to take advantage of possibilities like those listed above and more. In the *Research* section, we highlight how pre-service education still tends to compartmentalize and inadequately address the topic of technology. That section suggests a need to invest in research about how to deeply integrate technology in pre-service teacher training programs. In-service teachers, too, will need professional development to take advantage of opportunities that AI can provide, like those presented in the *Teaching* section. Professional development will need to be balanced not only to discuss opportunities but also to inform teachers of new risks, while providing them with tools to avoid the pitfalls of AI.

³⁹ Ersozlu, Z., Ledger, S., Ersozlu, A., Mayne, F., & Wildy, H. (2021). Mixed-reality learning environments in teacher education: An analysis of TeachLivETM Research. *SAGE Open*, 11(3). <https://doi.org/10.1177/21582440211032155>.

"Humans are well suited to discern the outcomes...because we are the ones that have the capacity for moral reflection and empathy. So, in other words, I want the AI to help me really quickly and easily see what my student needs in their learning journey."

—Sarah Hampton

By nature, teaching requires significant time in planning as well to account for the breadth of needs across their rosters—especially for inclusive learning environments and students with IEPs and 504 plans. AI could help teachers with recommendations that are tuned to their situation and their ways of practicing teaching and support with adapting found materials to fit their exact classroom needs. For students with an IEP, AI could help with finding components to add to lesson plans to fully address standards and expectations and to meet each student's unique requirements. Even beyond finding components, AI might help adapt standardized resources to better fit specific needs—for example, providing a voice assistant that allows a student with a visual difficulty to hear material and respond to it or permitting a group of students to present their project using American Sign Language (ASL) which could be audibly voiced for other students using an AI ASL-to-Spoken-English translation capability. Indeed, coordinating IEPs is time-consuming work that might benefit from supportive automation and customized interactivity that can be provided by AI.

Reflection is important too. In the bustle of a classroom, it is sometimes difficult to fully understand what a student is expressing or what situations lead to certain positive or negative behaviors. Again, context is paramount. In the moment, teachers may not be aware of external events that could shape their understanding of how students are showing up in their classrooms. Tools that notice patterns and suggest ways to share information might help students and teachers communicate more fully about strengths and needs.

Designing, Selecting, and Evaluating AI Tools

The broadest loop teachers should be part of is the loop that determines what classroom tools do and which tools are available. Today, teachers already play a role in designing and selecting technologies. Teachers can weigh in on usability and feasibility. Teachers examine evidence of efficacy and share their findings with other school leaders. Teachers already share insights on what is needed to implement technology well.

While these concerns will continue, AI will raise new concerns too. For example, the following *Formative Assessment* section raises concerns about bias and fairness that can lead to algorithmic discrimination. Those concerns go beyond data privacy and security; they raise attention to how technologies may unfairly direct or limit some students' opportunities to learn. A key takeaway here is that teachers will need time and support so they can stay abreast of both the well-known and the newer issues that are arising and so they can fully participate in design, selection, and evaluation processes that mitigate risks.

Challenge: Balancing Human and Computer Decision-Making

One major new challenge with AI-enabled tools for teachers is that AI can enable autonomous activity by a computer, and thus when a teacher delegates work to an AI-enabled tool, it may

carry on with that work somewhat independently. Professor Inge Molenaar⁴⁰ has wondered about the challenges of control in a hybrid teaching scenario: When should a teacher be in control? What can be delegated to a computational system? How can a teacher monitor the AI system and override its decisions or take back control as necessary?

Figure 8: The tension between human and AI decision making: Who is in control?




Figure 8 expresses the tension around control. To the left, the teacher is fully in control, and there is no use of AI in the classroom. To the right, the technology is fully in control with no teacher involved—a scenario which is rarely desirable. The middle ground is not one dimensional and involves many choices. Molenaar analyzed products and suggests some possibilities:

- The technology only offers information and recommendations to the teacher.
- The teacher delegates specific types of tasks to the technology, for example, giving feedback on a particular math assignment or sending out reminders to students before an assignment is due.
- The teacher delegates more broadly to the technology, with clear protocols for alerts, for monitoring, and for when the teacher takes back control.

These and other choices need to be debated openly. For example, we may want to define instructional decisions that have different kinds of consequences for a student and be very careful about delegating control over highly consequential decisions (for example, placement in a next course of study or disciplinary referrals). For human in the loop to become more fully realized, AI technologies must allow teacher monitoring, have protocols to signal a teacher when their judgment is needed, and allow for classroom, school, or district overrides when they disagree with an instructional choice for their students. We cannot forget that if a technology allows a teacher choice—which it should—it will take significant time for a teacher to think through and set up all the options, requiring greater time initially.

Challenge: Making Teaching Jobs Easier While Avoiding Surveillance


We also recognize that the very technologies that make jobs easier might also introduce new possibilities for surveillance (Figure 9). In a familiar example, when we enable a voice assistant in the kitchen, it might help us with simple household tasks like setting a cooking timer. And yet the same voice assistant might hear things that we intended to be private. This kind of dilemma will

⁴⁰ Molenaar, I. (2022). Towards hybrid human-AI learning technologies. *European Journal of Education*, 00, 1–14. <https://doi.org/10.1111/ejed.12527>

occur in classrooms and for teachers. When they enable an AI-assistant to capture data about what they say, what teaching resources they search for, or other behaviors, the data could be used to personalize resources and recommendations for the teacher. Yet the same data might also be used to monitor the teacher, and that monitoring might have consequences for the teacher. Achieving trustworthy AI that makes teachers' jobs better will be nearly impossible if teachers experience increased surveillance.

A related tension is that asking teachers to be “in the loop” could create more work for teachers if not done well, and thus, being in the loop might be in tension with making teaching jobs easier. Also related is the tension between not trusting AI enough (to obtain assistance) or trusting it too much (and incurring surveillance or loss of privacy). For example, researchers have documented that people will follow instructions from a robot during a simulated fire emergency even when (a) they are told the robot is broken and (b) the advice is obviously wrong.⁴¹ We anticipate teachers will need training and support to understand how and when they will need to exercise human judgement.

Figure 9: Highly customized assistance vs. increased teacher surveillance



Challenge: Responding to Students' Strengths While Protecting Their Privacy

Educators seek to tackle inequities in learning, no matter how they manifest locally (e.g. in access to educational opportunities, resources, or supports). In culturally responsive⁴² and culturally sustaining⁴³ approaches, educators design materials to build on the “assets”—individual, community, and cultural strengths that students bring to learning. Along with considering assets, of course, educators must meet students where they are, including both strengths and needs. AI could assist in this process by helping teachers with customizing curricular resources, for example. But to do so, the data inputted in an AI-enabled system would have to provide more information about the students. This information could be, but need not be, demographic details. It could also be information about students' preferences, outside interests, relationships,


⁴¹ Wagner, A.R., Borenstein, J. & Howard, A. (September 2018). Overtrust in the robotics age. *Communications of the ACM*, 61(9), 22-24. <https://doi.org/10.1145/3241365>

⁴² Gay, G. (2018). *Culturally responsive teaching: Theory, research, and practice*. Teachers College Press. ISBN: 978-0807758762

⁴³ Paris, D., & Alim, H.S. (Eds.). (2017). *Culturally sustaining pedagogies: Teaching and learning for justice in a changing world*. Teachers College Press. ISBN: 978-0807758342

or experiences.⁴⁴ What happens to this data, how it is deleted, and who sees it is of huge concern to educators. As educators contemplate using AI-enabled technologies to assist in tackling educational inequities, they must consider whether the information about students shared with or stored in an AI-enabled system is subject to federal or state privacy laws, such as FERPA. Further, educators must consider whether interactions between students and AI systems create records that must be protected by law, such as when a chatbot or automated tutor generates conversational or written guidance to a student. Decisions made by AI technologies, along with explanations of those decisions that are generated by algorithms may also be records that must be protected by law. Therein, a third tension emerges, between more fully representing students and protecting their privacy (Figure 10).

Figure 10: Responding to students' strengths while fully protecting student privacy



Further, representation would be just a start toward a solution. As discussed earlier in this report, AI can introduce algorithmic discrimination through bias in the data, code, or models within AI-enhanced edtech. Engineers develop the pattern detection in AI models using existing data, and the data they use may not be representative or may contain associations that run counter to policy goals. Further, engineers shape the automations that AI implements when it recognizes patterns, and the automations may not meet the needs of each student group with a diverse population. The developers of AI are typically less diverse than the populations they serve, and as a consequence, they may not anticipate the ways in which pattern detection and automation may harm a community, group, or individual.

AI could help teachers to customize and personalize materials for their students, leveraging the teacher's understanding of student needs and strengths. It is time consuming to customize curricular resources, and teachers are already exploring how AI chatbots can help them design additional resources for their students. An elementary school teacher could gain powerful supports for changing the visuals in a storybook to engage their students or for adapting language that poorly fits local manners of speaking or even for modifying plots to incorporate other dimensions of a teacher's lesson. In the *Learning* section, we noted that AI could help identify learner strengths. For example, a mathematics teacher may not be aware of ways in which a student is making great sense of graphs and tables about motions when they are in another teacher's physics classroom and might not realize that using similar graphs about

⁴⁴ Zacamy, J. & Roschelle, J. (2022). Navigating the tensions: How could equity-relevant research also be agile, open, and scalable? Digital Promise. <http://hdl.handle.net/20.500.12265/159>; Baker, R.S., Esbenshade, L., Vitale, J., & Karumbaiah, S. (2022). Using demographic data as predictor variables: A questionable choice. <https://doi.org/10.35542/osf.io/y4wvj>

motion could help with their linear function lesson. AI might help teachers when they seek to reflect student strengths by creating or adapting instructional resources.

Yet, the broad equity challenges of avoiding algorithmic discrimination while increasing community and cultural responsiveness must be approached within the four foundations we earlier outlined: human in the loop, equity, safety and effectiveness, and evaluation of AI models. We cannot expect AI models to respect cultural responsiveness. The Department is particularly concerned that equity is something that engaged educators and other responsive adults are in the best position to address and something that is never solely addressable as a computational problem.

Questions Worth Asking About AI for Teaching

As leaders in both pre-service and post-service teacher education contemplate how AI can improve teaching (along with policymakers, developers, and researchers), we urge all in the ecosystem to spend more time asking these questions:

- Is AI improving the quality of an educator's day-to-day work? Are teachers experiencing less burden and more ability to focus and effectively teach their students?
- As AI reduces one type of teaching burden, are we preventing new responsibilities or additional workloads being shifted and assigned to teachers in a manner that negates the potential benefits of AI?
- Is classroom AI use providing teachers with more detailed insights into their students and their strengths while protecting their privacy?
- Do teachers have oversight of AI systems used with their learners? Are they exercising control in the use of AI-enabled tools and systems appropriately or inappropriately yielding decision-making to these systems and tools?
- When AI systems are being used to support teachers or to enhance instruction, are the protections against surveillance adequate?
- To what extent are teachers able to exercise voice and decision-making to improve equity, reduce bias, and increase cultural responsiveness in the use of AI-enabled tools and systems?

Key Recommendation: Inspectable, Explainable, Overridable AI

In the Introduction, we discuss the notion that when AI is incorporated into a system, the core of the AI is a model. In the *Learning* section, we discuss that we need to be careful that models align to the learning we envision (e.g., that they aren't too narrow). Now, based on the needs of teachers (as well as students and their families/caregivers), we add another layer to our criteria for good AI models: the need for explainability.⁴⁵ Some AI models can recognize patterns in the world and do the right action, but they cannot explain why (e.g., how they arrived at the

⁴⁵ Khosravi, H., Shum, S.B., Chen, G., Conati, C., Tsai, Y.-S., Kay, J., Knight, S., Martinez-Maldonado, R., Sadiq, S., Gašević, D. (2022). Explainable artificial intelligence in education. *Computers and Education: Artificial Intelligence*, 3. <https://doi.org/10.1016/j.caeai.2022.100074>

connection between the pattern and the action). This lack of explainability will not suffice for teaching; teachers will need to know how an AI model analyzed the work of one of their students and why the AI model recommended a particular tutorial, resource, or next step to the student.

Thus, explainability of an AI system's decision is key to a teacher's ability to judge that automated decision. Such explainability helps teachers to develop appropriate levels of trust and distrust in AI, particularly to know where the AI model tends to make poor decisions. Explainability is also key to a teacher's ability to monitor when an AI system may be unfairly acting on the wrong information (and thus may be biased. We discuss bias and fairness more in the *Assessment* section next).

Surrounding the idea of explainability is the need for teachers to be able to inspect what an AI model is doing. For example, what kinds of instructional recommendations are being made and to which students? Which students are being assigned remedial work in a never ended loop? Which are making progress? Dashboards in current products present some of this information, but with AI, teachers may want to further explore which decisions are being made and for whom and know of the student-specific factors that an AI model had available (and possibly which factors were influential) when reaching a particular decision. For example, some of today's adaptive classroom products use limited recommendation models that only consider student success on the last three mathematics problems and do not consider other variables that a teacher would know to consider, such as whether a student has an IEP Plan or other needs.

Our call for attending to equity considerations as we evaluate AI models requires information about how discriminatory bias may arise in particular AI systems and what developers have done to address it. This can only be achieved with transparency for how the tools use datasets to achieve outcomes and what data they have available or that a teacher could include in her judgement but are not available to the system (IEP status is offered as an example above).

Teachers will also need the ability to view and make their own judgement about automated decisions, such as decisions about which set of mathematics problems a student should work on next. They need to be able to intervene and override decisions when they disagree with the logic behind an instructional recommendation.⁴⁶ Teachers need protection against adverse ramifications when they assert human judgement over an AI system's decision.

⁴⁶ Ruiz, P. & Fusco, J. (2022). *Teachers partnering with artificial intelligence: Augmentation and automation*. Digital Promise. <https://digitalpromise.org/2022/07/06/teachers-partnering-with-artificial-intelligence-augmentation-and-automation/>

"These systems sometimes are seen as a black box kind of a situation where predictions are made based on lots of data. But what we need is to have a clear view—to clearly show how those recommendations or those interactions are made and what evidence is used or what data is used to be able to make those recommendations so teachers and everyone involved know about why that kind of system is providing that type of information. So, having open learning environments or inspectable learner models or applications where the stakeholders can understand how these systems make decisions or recommendations is going to be an important aspect in the future of teaching and learning."

—Diego Zapata-Rivera

Formative Assessment

Formative assessment is traditionally a key use of edtech because feedback loops are vital to improving teaching and learning.⁴⁷ As we have emphasized throughout this report, a top priority with AI is to keep humans in the loop and in control, which includes focusing on the people engaged with formative assessments: students, teachers, school leaders, families/caregivers, and others who support learners. In the definition below, please note the overlap between definitions of AI and formative assessment; both have to do with detecting patterns and choosing a future course of action (that adapts to learner strengths and needs).

Assessment refers to all those activities undertaken by teachers, and by the students in assessing themselves, which provide information to be used as feedback to modify the teaching and learning activities in which they are engaged. Such assessment becomes "formative assessment" when the evidence is actually used to adapt the teaching to meet the needs.⁴⁸

Building on Best Practices

A number of dimensions hold potential for shaping the future of formative assessments, and many have ready extensions to the field of AI-enabled systems and tools. For example, the 2017 NETP discussed how technology can lead to improved formative assessments along seven dimensions, listed below:

- 1. Enabling Enhanced Question Types:**
to give students more ways to show what they know and can do.
- 2. Measurement of Complex Competencies:**
to better elicit growth in important skills that go beyond typical subject matter standards, for example, in measuring practices, social skills like teamwork, self-regulation, and work-relevant skills (e.g., making presentations or leading teams).
- 3. Providing Real-Time Feedback:**
to maintain and increase student engagement and to support effective learning, providing timely and helpful responses and suggestions to each learner.
- 4. Increasing Accessibility:**
to include neurodiverse learners and to engage learners' best communication capabilities as they share what they know and can do.

⁴⁷ Shute, V.J. (2008). Focus on formative feedback. *Review of Educational Research*, 78(1), 153–189. <https://doi.org/10.3102/0034654307313795>

⁴⁸ Black, P. & Wiliam, D. (1998). Inside the black box: Raising standards through classroom assessment. *Phi Delta Kappan*, 92(1), 81-90. <https://kappanonline.org/inside-the-black-box-raising-standards-through-classroom-assessment/>

5. **Adapting to Learner Ability and Knowledge:**
to make assessments more precise and efficient.
6. **Embedded Assessment in the Learning Process:**
to emphasize an assessment's role in improving teaching and learning (this report does not focus on assessment for accountability purposes).
7. **Assess for Ongoing Learning:**
to reveal progress over time and not just predetermined milestones.

AI models and AI-enabled systems may have potential to strengthen formative assessments. In one example, a question type that invites students to draw a graph or create a model can be analyzed with AI algorithms,⁴⁹ and similar student models might be grouped for the teacher to interpret. Enhanced formative assessment may enable teachers to better respond to students' understanding of a concept like "rate of change" in a complex, real-world situation. AI can also give learners feedback on complex skills, such as learning American Sign Language⁵⁰ or speaking a foreign language⁵¹ and in other practice situations where no person is available to provide immediate feedback.

Generally, an AI assistant may be able to reduce the load for teachers related to grading simpler aspects of student responses, allowing the teacher to focus their specialized judgment on important qualities of a whole essay or a complex project. We also may be able to better provide feedback with accessibility. For example, an AI-enabled learning technology may be able to interact verbally with a student about their response to an essay prompt, asking questions that guide the student to clarify their argument without requiring the student to read a screen or type at a keyboard. In the examples shared earlier in the *Learning* section, we also see that AI can be embedded in the learning process, providing feedback to students as they work to solve a problem, rather than only later after the student has reached a wrong answer. When formative assessment is more embedded, it can better support learning, and timely feedback is critical.⁵²

Although there are many points of connection like these between AI and formative assessments, our listening sessions also revealed attendees' desire to tackle some existing shortcomings in the field of formative assessment; namely, the time-consuming and sometime onerous nature of taking tests, quizzes, or other assessments and the lack of perceived value in the feedback loop by teachers and students.

Implications for Teaching and Learning

Real-time instructional feedback can be beneficial when it helps learners and teachers to improve. But common experience too often leaves students and teachers with unpleasant feelings toward assessment and thus poses a provocative conflict between the potential benefits

⁴⁹ Zhai, X., He, P., Krajcik, J. (2022). Applying machine learning to automatically assess scientific models. *Journal of Research in Science Teaching*. <https://doi.org/10.1002/tea.21773>

⁵⁰ Shao, Q., Sniffen, A., Blanchet, J., Hillis, M.E., Shi, X., Haris, T.K., & Balkcom, D. (2020). Teaching american sign language in mixed reality. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(4), 1-27. <https://doi.org/10.1145/3432211>

⁵¹ Godwin-Jones, R. (2021). Big data and language learning: Opportunities and challenges. *Language Learning & Technology*, 25(1), 4–19. <http://hdl.handle.net/10125/44747>

⁵² Wiggins, G. (2015). *Seven keys to effective feedback*. ACSD. <https://www.ascd.org/el/articles/seven-keys-to-effective-feedback>

of data collected through formative assessments and the practical implications of administering additional assessments in classrooms and schools.

Some AI-enabled systems and tools seek to address this potential conflict. For example, one AI-enabled reading tutor listens to students as they read aloud and provides on-the-spot feedback to improve their reading.⁵³ Students reportedly enjoyed reading aloud, and the approach was effective. Researchers have also embedded formative assessments in games so that students can show how well they understand Newtonian physics as they play increasingly difficult levels of a game.⁵⁴ If a student can more easily ask for and receive help when they feel frustrated or confused, reducing those feelings can feel encouraging. Student feelings of safety, confidence, and trust in the feedback generated by these AI-enabled systems and tools are essential to showcase their learning. That focus on learning growth and gains is optimal (absent negative consequences or a high-stakes environment).⁵⁵

AI-enhanced formative assessments may have the potential to save teachers' time (e.g., time spent on grading), allowing the instructor to spend more time engaged in helping students. AI-enhanced assessments may also benefit teachers if they provide detailed insights about student strengths or needs that may not be visible and if they support instructional adaptation or improvement by suggesting a small set of evidence-based recommendations for helping students master content. Such assessments may also be helpful outside of the classroom if it can provide feedback when the teacher is not available, for example, in completing homework or practicing a concept during study hall. As we discussed in the *Teaching* section, an essential aspect of deploying AI-based formative assessment must be centering teachers in system design.

Insight: AI Can Enhance Feedback Loops

The term “formative assessment” does not singularly mean a test or a measurement. Assessment becomes formative when it results in useful reflections and changes to the course of teaching, learning, or both.⁵⁶ The term “feedback loops” emphasizes that measurement is only part of the process. Feedback loops that lead to instructional improvement—including adaptations in teaching and learning—yield the strongest outcomes for students.

We also use “feedback loops” as a plural term because there are many types and levels of loops that are important. Students can benefit from feedback when they work individually, as a member of a small group, or in a classroom discussion. Feedback loops are valuable “in the moment”—for example, as a student practices a skill. Further, feedback loops are valuable when they cover larger spans of effort and reflections, such as at the end of presenting a project or term paper. In addition, feedback loops can assist teachers, for example, helping them notice

⁵³ Mostow, J., Aist, G., Burkhead, P., Corbett, A., Cuneo, A., Eitelman, S., Huang, C., Junker, B., Sklar, M.B., & Tobin, B. (2003). Evaluation of an automated reading tutor that listens: Comparison to human tutoring and classroom instruction. *Journal of Educational Computing Research*, 29(1), 61–117. <https://doi.org/10.2190/06AX-QW99-EQ5G-RDCF>

⁵⁴ Shute, V.J., Ventura, M., & Kim, Y.J. (2013). Assessment and learning of qualitative physics in Newton's Playground. *The Journal of Educational Research*, 106(6), 423–430. <https://doi.org/10.1080/00220671.2013.832970>

⁵⁵ Shute, V.J. (2008). Focus on formative feedback. *Review of Educational Research*, 78(1), 153–189. <https://doi.org/10.3102/0034654307313795>

⁵⁶ Black, P., & Wiliam, D. (2009). Developing the theory of formative assessment. *Educational Assessment, Evaluation and Accountability*, 21(1), 5–31. <https://doi.org/10.1007/s11092-008-9068-5>

their own patterns of responding to students' ideas. Moreover, feedback loops are critical to the continuous improvement of products and the implementation of programs.

Due to the importance of feedback loops, formative assessment could be a leading area for schools' explorations of powerful uses of AI in teaching and learning. Educators can build upon alignments between their long-standing visions for formative assessment and the emerging capabilities that AI holds. Further, the professional assessment community brings a toolkit for asking and answering questions about topics like bias and fairness. The psychometric toolkit of methods is a strong start toward the questions that must be asked and answered because it already contains ways to measure bias and fairness and, more generally, to benchmark the quality of formative assessments. But as our discussion reveals, AI can only make feedback loops better if we keep a firm eye on the weaknesses of AI and how AI introduces new concerns.

An Example: Automated Essay Scoring

One instructive example is Automated Essay Scoring (AES). To become strong writers, which is a valuable life skill, students need regular and specific feedback. However, reviewing and providing feedback on essays is very time consuming for humans. Hence, Ellis Page provided a first vision for computer programs that could review and provide feedback on student essays in 1966⁵⁷, and much effort has gone into AES technologies in the intervening 56 years. Many research review articles are available to summarize the progress, which has been impressive.⁵⁸ Further, some of today's applications of AES technologies will be familiar to readers, such as Grammarly, Turnitin, and the various essay analysis engines used by publishers and assessment companies. Also note that while the traditional AES functionality emphasizes scoring or rating essays, newer AI-enabled products focus more on providing students with constructive criticism and developing their skills as writers. Writing is a life skill that is important to the pursuit of college and career ambitions, and developing writers require comprehensive feedback. If developers could inexpensively augment human feedback to developing writers with AI feedback, it's possible that support for learning to write could become more equitable.

And yet, AES is an instructive example because researchers have analyzed limitations, too.⁵⁹ AES technologies in AI can analyze some features of student essays but can also be misled by the length of an essay, by a student who places appropriate keywords in sentences that don't make sense, and other flaws that a human reader would easily notice. In a telling quote, one team that reviewed the state of the art wrote this:

The authors further note that while human and AI judgements of essays may correlate, people and computers are not noticing the same things in student writing. Due to these limitations, we must continue to emphasize a human in the loop foundation for AI-enhanced formative assessment. AI may support but not replace high-quality, human-led processes and practices of formative assessment in schools.

⁵⁷ Page, E.B. (1966). The imminence of grading essays by computer. *Phi Delta Kappan*, 47(5), 238–243

⁵⁸ Ke, Z., & Ng, V. (2019). Automated essay scoring: A survey of the state of the art. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 6300–6308. <https://doi.org/10.24963/ijcai.2019/879>

⁵⁹ Doewes, A. & Pechenizkiy, M. (2021). On the limitations of human-computer agreement in automated essay scoring. In *Proceedings of the 14th International Conference on Educational Data Mining (EDM21)*. https://educationaldatamining.org/EDM2021/virtual/static/pdf/EDM21_paper_243.pdf

"Nevertheless, the time when AES systems will be able to operate on a par with human judges, with similar levels of connoisseurship for such features as meaning, emotion, originality, creativity, fluency, sense of audience and so on, arguably remains a long way off."

—Gardner, O'Leary, and Yuan⁶⁰

Key Opportunities for AI in Formative Assessment

Based on the listening sessions we held, we see three key areas of opportunity in supporting formative assessment using AI systems and models.

First, we recommend a strong focus on **measuring what matters**⁶¹ and particularly those things that have not been easily measured before and that many constituents would like to include in feedback loops. The example above, AES, was chosen because writing remains a valuable academic, workplace, and life skill. Looking at community goals through the lens of their visions for their high school graduates, we see that families/caregivers, students, and community leaders want to nurture graduates who solve problems adaptively, who communicate and collaborate well, who persevere and self-regulate when they experience challenges. "What matters" today reaches beyond a sole focus on the core academic content measured by large-scale summative assessments, to support students and teachers with actionable feedback that nurtures the broader skills students need to succeed and thrive. Further, within core academic content, AI may help us to provide feedback on the more realistic and complex aspects of doing math, for example, or investigating scientific phenomena, understanding history, or discussing literature.

Second, we'd like to see a strong focus on **improving help-seeking and help-giving**.⁶² Asking for and giving help is crucial to learning⁶³ and practicing a growth-mindset and central to the notion of human feedback loops. Students may not always know when they need help. In one example, computer algorithms can detect a student who is "wheel spinning" (working hard on mastering content but not making progress).⁶⁴ A student who is working hard may not feel like they need help, and the teacher may not be aware that the student is struggling if he or she appears to be "on task." AI may also be helpful by highlighting for students and teachers what forms of assistance have been most useful to the student in the recent past so that an educator can expand access to specific assistance that works for that individual student. Finally, educators may learn things from AI-enabled systems and tools that give feedback and hints during the completion of

⁶⁰ Gardner, J., O'Leary, M. & Yuan, L. (2021). Artificial intelligence in educational assessment: "Breakthrough? Or buncombe and ballyhoo?" *Journal of Computer Assisted Learning*, 37(5), 1207–1216. <https://doi.org/10.1111/jcal.12577>

⁶¹ Merrill, S. (2020). In schools, are we measuring what matters? *Edutopia*. <https://www.edutopia.org/article/schools-are-we-measuring-what-matters>

⁶² Roll, I., Aleven, V., McLaren, B.M., Koedinger, K.R. (2011). Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system, *Learning and Instruction*, 21(2), 267–280. <https://doi.org/10.1016/j.learninstruc.2010.07.004>.

⁶³ Webb, N.M., & Farivar, S. (1994). Promoting helping behavior in cooperative small groups in middle school mathematics. *American Educational Research Journal*, 31(2), 369–395. <https://doi.org/10.3102/00028312031002369>

⁶⁴ Kai, S., Almeda, M.V., Baker, R. S., Heffernan, C., & Heffernan, N. (2018). Decision tree modeling of wheel-spinning and productive persistence in skill builders. *Journal of Educational Data Mining*, 10(1), 36–71. <https://doi.org/10.5281/zenodo.3344810>

homework, utilizing that feedback to later reinforce concepts in direct instruction and strengthen the one-on-one support provided to students.⁶⁵ AI-enabled systems and tools can provide teachers with additional information about the students' recent work, so their instructor has a greater contextual sense as they begin to provide help.

Third, we advocate for **teachers and students to be strongly involved in designing feedback loops** as developers produce AI-enhanced formative assessments so they can directly voice what would make assessments less onerous and more convenient and valuable to them.⁶⁶ Earlier in the *Teaching* section, we emphasized how important it is to involve teachers in designing, selecting, and evaluating AI-enhanced technologies. Students need to be centered, too. They are experiencing AI in their everyday lives, and they have strong opinions on what is valuable and safe. There are local and cultural variations in how people provide and receive feedback, so adjusting feedback to align with community norms is important.

Key Recommendation: Harness Assessment Expertise to Reduce Bias

Bias and fairness are important issues in assessment design and administration,⁶⁷ and they hold relevance for the area of AI-enabled assessment. In traditional assessment, a test item might be biased if unnecessary details are included that differentially advantage some students (e.g., a story-based item that references a sport that only boys play regularly may be less helpful to girls). As discussed earlier, with AI, we now must worry about algorithmic discrimination which can arise due to the manner in which AI algorithms are developed and improved from large datasets of parameters and values that may not represent all cohorts of learners.

Algorithmic discrimination is not just about the measurement side of formative assessment; it is also about the feedback loop and the instructional interventions and supports that may be undertaken in response to data collected by formative assessments. There is a question both about access to such interventions and the quality or appropriateness of such interventions or supports. When an algorithm suggests hints, next steps, or resources to a student, we have to check whether the help-giving is unfair because one group systematically does not get useful help which is discriminatory. Fairness goes beyond bias as well. In AI-enabled formative assessment, both the opportunity to learn through feedback loops, as well as the quality of learning in and outside of such loops, should be addressed. Issues of bias and fairness have arisen in traditional assessments, and the field of psychometrics has already developed valuable tools to challenge and address these issues.⁶⁸ Assessment as a field may have a head start on tackling bias and fairness for AI in education. And yet the issues expand with AI, so the work is not done. Strong and deliberate attention to bias and fairness is needed as future formative assessments are developed.

⁶⁵ Walker, E., Rummel, N. & Koedinger, K.R. (2015). Adaptive intelligent support to improve peer tutoring in algebra. *International Journal of Artificial Intelligence in Education*, 24, 33–61 <https://doi.org/10.1007/s40593-013-0001-9>

⁶⁶ Swiecki, Z., Khosravi, H., Chen, G., Martinez-Maldonado, R., Lodge, J.M., Milligan, S., Selwyn, B. & Gašević, D. (2022). Assessment in the age of artificial intelligence. *Computers and Education: Artificial Intelligence*, 3. k <https://doi.org/10.1016/j.caeai.2022.100075>

⁶⁷ Reynolds, C.R., & Suzuki, L.A. (2012). Bias in psychological assessment: An empirical review and recommendations. *Handbook of Psychology, Second Edition*. <https://doi.org/10.1002/9781118133880.hop210004>

⁶⁸ Kaplan, R.M., & Saccuzzo, D.P. (2017). *Psychological testing: Principles, applications, and issues*. Cengage Learning.

Related Questions

As indicated, formative assessment is an area in which AI is expanding along a continuum that can be guided by visions already in place, such as the 2017 NETP. It is an area in which AI is poised to grow, especially with capabilities that power more feedback loops in student learning. As this growth takes place, we suggest ongoing attention to the following questions:

- Is formative assessment bringing benefits to the student learning experience and to the efficacy of classroom instruction?
- Are humans being centered in AI-enabled formative assessment and feedback loops?
- Are we providing empowering professional development to teachers so they can leverage feedback loops and safeguard against concerns?
- To what extent are the developers and implementers of AI-enabled systems and tools tackling new sources of algorithmic bias and continuing to make assessment fairer?
- Are governance policies regarding who owns, controls, and can view or use AI-enabled formative assessment data appropriate and adequate?
- Do we have sufficient guardrails against misuse of formative assessment data or automatically generated interpretations of student achievement and learning, such as on dashboards?
- Is trust in an AI-enabled assessment system, feedback loops, and data generated by such assessments growing or diminishing?

Research and Development

Policy relies upon research-based knowledge; likewise, improving practice depends on feedback loops that analyze empirical evidence. Consequently, the 2010 NETP specified a series of “grand challenges” which were “R&D problems that might be funded and coordinated at a national level.” One 2010 NETP grand challenge was to create personalized learning systems that continuously improve as they are used:

“Design and validate an integrated system that provides real-time access to learning experiences tuned to the levels of difficulty and assistance that optimize learning for all learners and that incorporates self-improving features that enable it to become increasingly effective through interaction with learners.”⁶⁹

Since 2010, much R&D has addressed this challenge. Conferences about learning analytics, educational data mining, and learning at scale have blossomed. Developers have created platforms that use algorithms and the analysis of big data to tune learning experiences. The challenge has not been fully achieved, and further work on this challenge is still relevant today.

Insight: Research Can Strengthen the Role of Context in AI

Despite the relevance of 2010’s grand challenges, it has become apparent that the R&D community is now looking to expand their attention. The 2010 challenges were stated as technical problems. Today’s researchers want to more deeply investigate context, and today’s tech companies want to develop platforms that are responsive to the learners’ characteristics and situations more broadly—not just in terms of narrow cognitive attributes. We see a push to transform R&D to address context sensitivity. We look forward to new meanings of “adaptive” that broaden outward from what the term has meant in the past decade. For example, “adaptive” should not always be a synonym of “individualized” because people are social learners. Researchers therefore are broadening “adaptivity” to include support for what students do as they learn in groups, a form of learning that is prevalent in schools across the U.S.

The focus on context is not an accident. Context is a traditional challenge in AI.⁷⁰ Thus, researchers and developers are wise to prioritizing context. Unless we invest more in AI that is context-sensitive, it is quite likely that AI will break and fail to achieve educational goals.


Agreeing to prioritize context won’t be easy. As illustrated above in Figure 12, there will be a tension between depth of context and pace of technological advances in AI R&D. On the one hand, AI is sometimes presented as a race to be the first to advance new techniques or scale new applications—innovation is sometimes portrayed as rapidly going to scale with a minimally viable product, failing fast, and only after failure, dealing with context. On the other hand, researchers and developers see that achieving good innovations with AI in education will clearly

⁶⁹ U.S. Department of Education, Office of Educational Technology. (2010). *Transforming American Education: Learning Powered by Technology*. U.S. Department of Education. p. 78

⁷⁰ Boden, M.A. (2018). *Artificial intelligence: A very short introduction*. Oxford. ISBN: 978-0199602919

require bringing more context into the process early and often. For example, researchers highlight that humans must be continually adjusting the goals for technology and have noted that when we set forth goals, we often don't yet fully understand context; and as we learn about context, the goals must change.⁷¹ This suggests that context must be prioritized early and habitually in R&D; we don't want to win a race to the wrong finish line.

Figure 12: The tension between depth of context and pace of technological advances in AI



Further, intensifying focus on context in this work will change the nature of the R&D. There won't be just one type of change in R&D because context has multiple meanings. Attendees in our listening sessions described four types of context necessary for the future.

We list these four types of context below and then expand on each one in its own section. These four types emerged as topics of provocations to think differently about R&D but certainly do not exhaust the important ways of investigating context.

1. **Focus on the Long Tail:** How could we use big data and AI to pay more attention to the “long tail” of edtech use—going beyond a few “most typical” ways of using emerging technology and instead solving for digital equity and inclusion?
2. **Partnership in Design-Based Research:** How can we change who is involved and influential in designing the future of AI in education to more centrally include students, teachers, and other educational constituents?
3. **Connect with Public Policy:** How can work on AI in education build on general advances in AI ethics, safety, and regulation and contribute additional advances specific to educational policy?
4. **Rethink Teacher Professional Development:** How can we solve for new systems of teacher professional development (both pre-service and in-service) that align to the increasingly core role of technology in the teaching profession?

⁷¹ Russell, S. (2019). Human compatible: Artificial intelligence and the problem of control. Penguin. ISBN: 9780525558637


"We can't necessarily always apply traditional research methodologies to this topic because educational technology changes so quickly."

—Kristina Ishmael, Office of Educational Technology

Attention to the Long Tail of Learner Variability

At the core of R&D of AI in education, innovators will be building models that fit available data. The increasing scale and prevalence of technologies means that the data is coming from and including a wide range of different contexts and varied ways that people in those contexts engage in teaching and learning. Researchers in our listening sessions drew attention to the promise of AI for addressing “context” by reference to the long tail of learner variability.

Figure 13: The long tail of learner variability



As depicted in Figure 13, learners vary in their strengths and needs. The most frequently occurring mix of strength and needs (also known as “teaching to the middle”) is depicted leftmost, with less frequently occurring mixes spreading to the right. Rising upward, the figure depicts the number of learners who benefit from a particular learning design, pathway, or approach. We argue that AI can bring opportunities to address a wider spectrum of strengths and needs but only if developers and innovators focus on the long tail and not only “teaching to the middle.”

For the sake of argument, the figure indicates three zones. In a first zone, curricular resources are mostly standardized, with perhaps a dimension or two of adaptivity. For example, many existing products adapt based on the correctness of student answers and may also provide options to read or hear text in a second language. However, the core of the instructional approach is highly standardized. In a second zone, there is greater balance between how much standardization and how much adaptivity students can access. [Universal Design for Learning \(UDL\)](#) is one set of recommendations for providing learning opportunities in multiple formats and for

accommodating different learning progressions.⁷² UDL can enable accommodating more ways in which learners vary, and as teachers know, there are many more important ways to adapt to students than found in today's edtech products.

Students are neurodiverse. They bring different assets from their experiences at home, in their communities, and in their cultures. They have different interests and motivations. And they learn in varied settings—classrooms and schools differ, and at-home students learn in informal settings in ways that could complement school learning. These are all important dimensions of “context.” Zone 3 indicates highly adaptive learning, where standardization is less successful and where we need to discover a wider variety of approaches to engage learners and sustain powerful learning. Researchers in our listening sessions noted the promise of Zone 3 because AI's ability to recognize patterns in data can extend beyond the most common patterns and because AI's ability to generate customized content can extend beyond what people can reasonably generate on their own.

Notice that although the Zone 1 bar appears to be the tallest, and thus tends to attract initial attention, there are more students in Zones 2 and 3, the regions where AI can provide more help. Thus, it's important to ask where AI researchers and developers are directing their attention. When we say a model “fits,” are we saying it fits the most common and typical uses by teachers and learners? This sort of R&D is easier to do. However, machine learning and AI also can tailor a model to the less common and more culturally specific contexts, too. Therefore, how can constituents cultivate interdisciplinary expertise to direct attention among researchers and developers to focus on the long tail? If we do, the quality of what we do for those represented in that tail can be more adaptive and more context-sensitive. And to be most effective, it will require the integration of contextual, content, and technical expertise.

Within the long-tail challenge, the community is wondering how we can get to research insights that are both general and specific enough. When research produces very general abstractions about learning, it often doesn't give developers enough guidance on exactly how to adjust their learning environments. Conversely, when research produces a specific adaptive algorithm that works on one educational platform, it often remains hard to apply to additional platforms; research can be too detailed as well. The research community is also thinking about new partnerships that could bring more data and more diverse perspectives to the table, the topic of the next section.

Focusing on the long tail of learner variability is particularly important to addressing a long-standing key research question: “*Do new AI-enhanced approaches work to improve learning, and for whom and under what conditions?*”

Partnership in Design-Based Research

Of course, teachers must be included in rethinking their own professional development. This thought leads to another priority aspect of context: partnership in design-based research. With regard to inclusive design, attendees in our listening sessions brought up a variety of co-design⁷³

⁷² Rose, D. (2000). Universal design for learning. *Journal of Special Education Technology*, 15(4), 47-51. <https://doi.org/10.1177/016264340001500407>

⁷³ Roschelle, J., Penuel, W., & Shechtman, N. (2006). Co-design of innovations with teachers: definition and dynamics. In *Proceedings of the 7th International Conference on Learning Sciences*, Bloomington, IN. <https://doi.org/10.22318/icsl2006.606>

and other participatory processes and goals that can be used in R&D.⁷⁴ By co-design, they mean sharing power with non-researchers and non-developers through all the phases of design and development, which would result in more influence by teachers, students, and other constituents in the shape of AI-enabled edtech. The shift toward co-design was palpable throughout our listening sessions, but as researchers and developers have not standardized on one particular co-design method, we share some representative examples.

- Youth can powerfully participate in design when researcher methods include participant co-design. Such research can investigate how to improve edtech while educating students. A listening session attendee asked about developing students' awareness of what data are being collected and how data are being used by developers.
- There is a near future need to go beyond representation so that co-designed solutions consider more generous contexts for broader possibilities, according to attendees.
- The shift of power dynamics is another research-worthy interest of the panel and attendees to understand the balance between a teacher's agency and a machine's suggestions.
- Likewise, such longitudinal research will require both the infrastructure and institutional support to fund necessary experimentation and requisite failures to elicit positive results and safe innovation.
- There is a desire for rapid cycle evaluations with inclusive feedback loops that return to the educators themselves as essential relative to traditional research approaches.
- Many researchers also mentioned a focus on explainable AI as essential to enable participation in the design and evaluation of emerging AI approaches in education.

The conversations raised this question: how can co-design provide an empowering form of participation in design and thus achieve digital inclusion goals? Such digital inclusion can span many layers of design, including diverse representation in design of policies around data, design of adaptivity, and other user experiences in AI systems, design of plans for cultivating AI literacy for users of new platforms, and lastly, the design of plans to evaluate systems.

Re-thinking Teacher Professional Development

With regard to teachers as professionals, both researchers and other educators attending our listening sessions were highly concerned about the disconnect between how teachers are prepared versus how they are expected to work with emerging technology. When we discuss learning, teachers are central actors, and thus the contexts in which they are prepared is centrally important to their ability to do great work in current and emerging technological environments.

Teacher professional development, professional learning, and leadership (PD or PL) for emerging technologies was seen as an area needing intense re-thinking, and research could lead the way. Today, few who prepare to become a teacher in an established pre-service program learn about the effective use of educational technology in schools and classrooms; those who do

⁷⁴ Center for Integrative Research in Computing and Learning Sciences (CIRCLS). (2022, Feb.). From Broadening to Empowering: Reflecting on the CIRCLS'21 Convening. <https://circls.org/circls21report>

have the opportunity to investigate technology rarely think about the structures that shape its use in the classroom and in educational leadership. Consequently, a troubling dichotomy arises between a small set of investigators who specifically consider educational technology in their research on teaching and a broader group of educators who see educational technology as a generic instructional resource. The challenge is high because teacher professional development will remain highly varied by local contexts. Yet insufficient attention to teachers as leaders in the use and further development of effective educational technology is widespread in teacher professional development research.

One response can be in terms of investigating how to nurture greater AI literacy for all teachers. AI literacy is not only important to protect educators and students from possible dangers but also valuable to support teachers to harness the good and do so in innovative ways. A panelist reminded the group that this work implies how we prepare educators with a baseline AI literacy and understanding. More transparency and authentic dialogue can foster trust, which was mentioned by a researcher as a chief concern for all teachers and students.

This is not to suggest that AI literacy is a complete or even a simple fix. Researchers want to ask fundamental questions about what it means for teachers to be professionals, especially as emerging technologies gain ground in schools and classrooms—our teachers' professional workplaces. Researchers want to broadly reconceptualize teacher professionalism and to stop treating technology as an add-on element of professional development.

Connecting with Public Policy

Defining human-centered AI for education requires the embrace of a human-centered principle and foundation for developing and formulating policies that govern the application and use of AI more generally throughout society. For example, power dynamics that arise between companies and consumers in society around issues like data ownership will also arise in the education-specific ecosystem. Further, the public discourse in which people are discussing ethics, bias, responsibility, and many other necessary concepts will be happening simultaneously in public policy and in educational ecosystems.

One clear implication in our listening sessions was that efforts to improve AI literacy in education could be important and helpful to society more generally. For example, one panelist said that an overarching goal of improving AI literacy is necessary if they are to contribute to how those technologies are designed. Another researcher was interested in how edtech can provide environments where students can experience having difficult discussions across perspectives, an issue which is endemic to present society. A third researcher noted the insufficiencies of prior efforts to contend with algorithmic bias, ethics, and inclusion due to a classroom's complex social dynamics.

Researchers want to take a lead in going beyond checkbox approaches to take these issues seriously. And they also acknowledge that engaging with policy is often a new form of context for edtech and AI researchers, many of whom don't have long experiences in policy arenas. Likewise, developers often do have experience with some policy issues, such as data privacy and security, but are now needing to become part of new conversations about ethics, bias,

transparency, and more, a problem that the EdSAFE AI Alliance is addressing through multi-sector working groups and policy advocacy.⁷⁵

Key Recommendation: Focus R&D on Addressing Context

Attendees who have participated in listening sessions leading up to this report were exceptionally clear that their view of future R&D involved a shift from narrow technical questions to richer contextual questions. This expansive shift toward context, as detailed below, is the foundational orientation that the listening session attendees saw as being necessary to advancing R&D.

Attendees included these as dimensions of context:

- learner variability, e.g., in disabilities, languages spoken, and other relevant characteristics;
- interactions with peers, teachers, and others in the learning settings;
- relationships across home, school, and community settings, including cultural assets;
- instructional resources available while learning;
- teacher preparation; and
- policies and systems that structure teaching and learning.

To more fully represent the context of teaching and learning, including these and other dimensions of text, researchers will have to work in partnership with others to understand which aspects of context are most relevant to teaching and learning and how they can be usefully incorporated into AI models.

Ongoing Questions for Researchers

As mentioned earlier, people are good at context; AI—not so much. R&D investment in context-rich edtech thus could serve multiple national interests because finding ways to do a better job with context would be a fundamental advancement in AI. Indeed, questions like these reverberate across all applications of AI in society, and education is a centrally good context for investigating them:

- Are AI systems moving beyond the tall portions of the “long tail” to adapt to a greater range of conditions, factors, and variations in how people learn?
- To what extent are AI technologies enhancing rather than replacing human control and judgment of student learning?
- How will users understand the legal and ethical implications of sharing data with AI enabled technologies and how to mitigate privacy risks?
- To what extent does technology account for the complex social dynamics of how people work and learn together, or is technology leading humans to narrow or oversimplify?
- How can we more clearly define what we mean by a context-sensitive technology in terms that are both concrete and broad enough? How can we measure it?

⁷⁵ Nentrup, E. (2022). How Policymakers Can Support Educators and Technology Vendors Towards SAFE AI. EdSAFE AI Alliance. <https://www.edsafeai.org/post/how-policymakers-can-support-ai-ed>

- To what extent are technical indicators and human observations of bias or unfairness working together with human observations? How can concerns about ethics and equity in AI technologies become actionable both in R&D, and later, when AI is widely used?
- Are we learning for whom and under what conditions AI systems produce desired benefits and impacts and avoid undesirable discrimination, bias, or negative outcomes?

Desired National R&D Objectives

Attendees sought immediate progress on some key R&D issues, such as these:

- Clarifying and achieving a consensus on the terms that go beyond data privacy and data security, including ideas like human-centered, value-sensitive, responsible, ethical, and safe so constituents can advocate for their needs meaningfully and consistently
- Creating and studying effective programs for AI literacy for students, teachers, and educational constituents in general, including literacy with regard to the ethics and equity issues specific to AI in educational settings
- Advancing research and development to increase fairness, accountability, transparency, and safety in AI systems used in educational settings
- Defining participatory or co-designed research processes that include educators in the development and conduct of research related to the development, use, and efficacy of AI-enabled systems and tools
- Highlighting and advancing R&D efforts that empower the participation and voices of youth regarding research, data, and design of AI applications for teaching and learning

Longer term desires for a national R&D program include some of the following objectives:

- Funding sustainable partnerships that uncover what context means and how it can be addressed over longer periods of time
- Better connecting goals for “broadening participation” (for example, in STEM learning pathways) to strategies for addressing learner variability and diversity
- Prioritizing research to revitalize support for instructors in light of the increasingly technological nature of K-12, higher education, and workplace learning settings
- Creating infrastructure and new ways of working together beyond individual field-initiated grants so that R&D with big data and leveraging emerging AI capabilities becomes safer and more productive

Recommendations

Earlier, we asked two guiding questions:

1. What is our collective vision of a desirable and achievable educational system that leverages automation while protecting and centering human agency?
2. On what timeline will we be ready with necessary guidelines and guardrails along with convincing evidence of positive impacts, so that we can ethically and equitably implement this vision widely?

Answers to the first question are provided throughout the *Learning, Teaching, Assessment, and Research* sections. This section turns to a call to action to education leaders and to recommendations. Core to the Department's perspective is that education will need leadership specific to our sector. Leadership should recognize and build on prior accomplishments in edtech (such as strong prior work on student privacy and school data security) as well as broad frameworks for safe AI (such as the *Blueprint for an AI Bill of Rights*). Leadership must also reach beyond these accomplishments and frameworks to address emerging opportunities and risks that are specific to novel capabilities and uses of AI in education.

Insight: Aligning AI to Policy Objectives

Individual sections of this policy report provided insights in each of four areas—learning, teaching, assessment, and research. These insights, synthesized from extensive stakeholder consultation and listening sessions, show that the advances in AI can bring opportunities to advance the Department's policy objectives:

- In support of our objective of attracting and retaining teachers, our nation could focus on AI assistants that make teaching jobs better and provide teachers with the information they need to work closely and empathically with students. An emphasis on teachers in the loop could ensure that AI-enabled classroom technologies keep teachers in the know, in touch with their students, and in control of important instructional decisions. Keeping the teacher in the loop is important to managing risks, as well.
- In support of equitable learning, especially for those most affected by the pandemic, AI could shift edtech from a current deficit-based model to a strengths-based alternative. In addition to finding student weaknesses and assigning fixes, edtech could make recommendations based on strengths that students bring to learning and how adapting to the whole student—a cognitive, social, and self-regulating person—could enable more powerful learning. Adapting to the whole student should include supporting students with disabilities as well as English learners. With regard to equity, we must remain highly attuned to the challenges of bias (which are inherent to how AI systems are developed) and take firm action to ensure fairness.
- With regard to growth trajectories to successful careers, AI-enabled assessments could provide students and teachers with formative guidance on a wider range of valuable skills, focusing on providing information that enhances learning. Aligned with the human-centric view, we should take a systems view of assessments where students, teachers, and others remain at the center of instructional decision making.

- With regard to equity, as research advances and brings more context into AI, we will be better able to use AI to support goals that require customization of learning resources, such as enabling teachers to more easily transform materials to support neurodiverse learners and increase responsiveness to local communities and cultures.

Going forward, educational leaders need to bring these and their own policy priorities to the table at every discussion about AI, driving the conversation around human priorities and not only their excitement about what new technology might do. Fundamentally, AI seeks to automate processes that achieve goals, and yet, AI should never set goals. The goals must come from educators' vision of teaching and learning and educators' understanding of students' strengths and needs.

Calling Education Leaders to Action

We summarize seven recommendations for policy action. These recommendations are for education leaders. In the introduction, we note the necessity of involving education constituents in determining policies for AI. We also observed throughout our listening sessions that people coming from many different roles in education all have passion, knowledge, and insights to contribute. In our view, all types of constituents can be education leaders. We are reluctant to suggest any constituent role is more important to advance any of the recommendations, but we call out specific needs for action within some of the recommendations where it is warranted.

Recommendation #1: Emphasize Humans in the Loop

We start with a central recommendation throughout this report. This recommendation was a clear constituent favorite. Indeed, across more than 700 attendees in our listening sessions, the predominant discussion tackled how constituents can achieve a consensus vision for AI-enabled edtech where humans are firmly at the center. The *Blueprint for an AI Bill of Rights* similarly calls for “access to timely human consideration and remedy by a fallback and escalation process if an automated system fails, it produces an error, or you would like to appeal or contest its impacts...” Building on this consensus, we call upon all constituents to adopt “humans in the loop” as a key criterion for educational use of AI.

We envision a technology-enhanced future more like an electric bike and less like robot vacuums. On an electric bike, the human is fully aware and fully in control, but their burden is less, and their effort is multiplied by a complementary technological enhancement. Robot vacuums do their job, freeing the human from involvement or oversight.

Although teachers should not be the only humans involved in loops, Figure 5 provided examples of three types of teacher loops that are central to education and can be used to illustrate what “human in the loop” means. Here, we use the example of an AI chatbot to elaborate on the meaning of the loops. First, as students become involved in extended interactions with AI chatbots, teachers will need to educate students about safe AI use, monitor their use, and provide human recourse when things go astray. Second, teachers are beginning to use chatbots to plan personalized instruction for their students; they will need to be involved in loops with other teachers to understand effective prompts, to know how to analyze AI-generated lesson plans for flaws, and to avoid the human tendency to overly trust AI systems and underapply human judgement. Third, teachers need to be involved in the design and evaluation of AI systems before they are used in classrooms and when needs for improvement are observed. In one example, to design AI-generated homework support for students, teachers' in-depth understanding of the

cognitive, motivational, and social supports their students need will provide much-needed guidance as a homework-support chatbot is designed.

In framing AI in education, this report advances a key recommendation of “human in the loop” AI because the phrase readily communicates a criterion that everyone can use as they determine which AI-enabled systems and tools are appropriate for use in teaching and learning. In a rather technical field, human in the loop is an approachable and humanistic criterion. Rather than suggesting that AI-enabled systems and tools should replace teachers, this term instead solidifies the central role of educators as instructors and instructional decision makers, while reinforcing the responsibility of teachers to exercise judgement and control over the use of AI in education. It resonates with the important idea of feedback loops, which are highly important to how people teach and learn. It also aligns with the ideas of inspectable, explainable, severable, and overridable AI.

The Department agrees with listening session participants who argued that teachers should not be the only humans in the loop and calls upon parents, families, students, policy makers, and system leaders to likewise examine the “loops” for which they are responsible, critically analyze the increasing role of AI in those loops, and determine what they need to do to retain support for the primacy of human judgement in educational systems.

Recommendation #2: Align AI Models to a Shared Vision for Education

*“All models are wrong, but some are useful.”
—George Box, Statistician*

As we have discussed across every section of this report, AI technologies are grounded in models, and these models are inevitably incomplete in some way. It is up to humans to name educational goals and measure the degree to which models fit and are useful—or don’t fit and might be harmful. Such an assessment of how well certain tools serve educational priorities may seem obvious, but the romance of technology can lead to a “let’s see what the tech can do” attitude, which can weaken the focus on goals and cause us to adopt models that fit our priorities poorly.

Here we call upon educational policy and decision makers at the local, state, and federal level to use their power to align priorities, educational strategies, and technology adoption decisions to place the educational needs of students ahead of the excitement about emerging AI capabilities. We want to strengthen their attention to existing state, district, and school-level policies that guide edtech adoption and use, such as the four levels of evidence in ESSA, the privacy requirements of FERPA, and enhanced policies to come. Local education leaders know best what their urgent educational priorities are. Every conversation about AI (or any emerging technology) should start with the educational needs and priorities of students front and center and conclude with a discussion about the evaluation of effectiveness re-centered on those needs and priorities. Equity, of course, is one of those priorities that requires constant attention, especially given the worrisome consequences of potentially biased AI models.

We especially call upon leaders to avoid romancing the magic of AI or only focusing on promising applications or outcomes, but instead to interrogate with a critical eye how AI-enabled systems and tools function in the educational environment. We ask leaders to distrust broad claims and ask six types of questions, listed below. Throughout this report, we elaborated on

which characteristics of AI model use in education are most important to evaluate for alignment to intended educational goals. To aid leaders, we summarize our insights about AI models and their use in educational tools and systems in Figure 14.

Figure 14: Recommendation for desired qualities of AI tools and systems in education



In this figure, we center teaching and learning in all considerations about the suitability of an AI model for an educational use. Humans remain in the loop of defining, refining, and using AI models. We highlight the six desirable characteristics of AI models for education (elaborating from principles in the *Blueprint for an AI Bill of Rights* to fit the specifics of educational systems):

1. **Alignment of the AI Model to Educators' Vision for Learning:** When choosing to use AI in educational systems, decision makers prioritize educational goals, the fit to all we know about how people learn, and alignment to evidence-based best practices in education.
2. **Data Privacy:** Ensuring security and privacy of student, teacher, and other human data in AI systems is essential.
3. **Notice and Explanation:** Educators can inspect edtech to determine whether and how AI is being incorporated within edtech systems. Educators' push for AI models can explain the basis for detecting patterns and/or for making recommendations, and people retain control over these suggestions.
4. **Algorithmic Discrimination Protections:** Developers and implementers of AI in education take strong steps to minimizing bias and promoting fairness in AI models.

5. **Safe and Effective Systems:** The use of AI models in education is based on evidence of efficacy (using standards already established in education for this purpose) and work for diverse learners and in varied educational settings.
6. **Human Alternatives, Consideration and Feedback:** AI models that support transparent, accountable, and responsible use of AI in education by involving humans in the loop to ensure that educational values and principles are prioritized.

Although we first address our recommendation to interrogate how educational systems use AI models to educational leaders who adopt technologies, other leaders also have integral roles to play. Teachers and students, as well as their families/caregivers, contribute significantly to adoption decisions also. And leaders and parents must support educators when they question or override an AI model based on their professional wisdom. Developers of technologies need to be forthcoming about the models they use, and we may need policymakers to create requirements for disclosure so that the marketplace can function on the basis of information about AI models and not only by the claims of their benefits.

We also emphasize the need for a government role. AI models are made by people and are only an approximation to reality. Thus, we need policies that require transparency about the AI models that are embedded in educational systems, as well as models that are inspectable, explainable, and overridable. Our listening sessions featured constituent calls for government doing more to hold developers accountable for disclosing the types of AI models they employ in large-scale products and the safeguards included in their systems. Government leaders can make a positive contribution to market conditions that enable building trust as AI systems are procured and implemented in education. We discuss these guidelines more in recommendation #4, which is about building trust.

Recommendation #3: Design Using Modern Learning Principles

We call for the R&D sector to ensure that product designs are based on best and most current principles of teaching and learning. The first decade of adaptivity in edtech drew upon many important principles, for example, around how to sequence learning experiences and how to give students feedback. And yet the underlying conception was often deficit-based. The system focused on what was wrong with the student and chose pre-existing learning resources that might fix that weakness. Going forward, we must harness AI's ability to sense and build upon learner strengths. Likewise, the past decade of approaches was individualistic, and yet we know that humans are fundamentally social and that learning is powerfully social. Going forward, we must build on AI capabilities that connect with principles of collaborative and social learning and which respect the student not just for their cognition but also for the whole human skill set. Going forward, we also must seek to create AI systems that are culturally responsive and culturally sustaining, leveraging the growth of published techniques for doing so. Further, most early AI systems had few specific supports for students with disabilities and English learners. Going forward, we must ensure that AI-enabled learning resources are intentionally inclusive of these students. The field has yet to develop edtech that builds upon each student's ability to make choices and to self-regulate in increasingly complex environments. We have to develop edtech that expands students' abilities to learn in creative modes and to expand their ability to discuss, write, present, and lead.

We also call upon educators to reject uses of AI that are based solely on machine learning from data—without triangulation based on learning theory and knowledge from practice. Achieving

effective and equitable educational systems requires more than processing “big data,” and although we want to harness insights from data, human interpretation of data remains highly important. We reject a technological determinism in which patterns in data, on their own, tell us what to do. Applications of AI in education must be grounded in established, modern learning principles, the wisdom of educational practitioners, and should leverage the expertise in the educational assessment community around detecting bias and improving fairness.

Recommendation #4: Prioritize Strengthening Trust

Technology can only help us to achieve educational objectives when we trust it. Yet, our listening sessions revealed the ways in which distrust of edtech and AI is commonplace. Constituents distrust emerging technologies for multiple reasons. They may have experienced privacy violations. The user experience may be more burdensome than anticipated. Promised increases in student learning may not be backed by efficacy research. They may have experienced unanticipated consequences. Unexpected costs may arise. Constituents may distrust complexity. Trust needs to incorporate safety, usability, and efficacy.


The Department firmly takes the stance that constituents want AI that supports teachers and rejects AI visions that replace teachers. And yet, teachers, students, and their families/caregivers need support to build appropriate levels of trust in systems that affect their work. In the broader ecosystem, trustworthy AI is recognized as a multidimensional problem (including the dimensions of Figure 14, above). If every step forward does not include strong elements of trust building, we worry that distrust will distract from innovation serving the public good that AI could help realize.

We expect that associations and societies have a key role in strengthening trust. Some important associations like the State Educational Technology Directors Association and the Consortium for School Network work with edtech leaders, and parallel organizations like EDUCAUSE work with postsecondary leaders. Other associations and societies work with teachers, education leaders, and education staff developers. Industry networks, like the EdSAFE AI Alliance, can bring together industry leaders to work together to foster trust. Additional societies bring researchers together. These societies and associations have the reach necessary to bring all parts of the educational ecosystem into discussions about trust and also the ability to represent the views of their constituents in cross-cutting policy discussions.

Recommendation #5: Inform and Involve Educators

Our listening sessions also asked for more specific direction on the question of what education leaders should do (see Figure 15). The most frequent responses fit three clusters: the need for guidelines and guardrails, strengthening the role of teachers, and re-focusing research and development. These are activities that constituents are asking for and that could expand trust. The recommendations that follow respond to these requests.

Figure 15: Listening session attendees prioritized involving practitioners, research, and evaluation and the need for guidelines and guardrails.



In particular, one concern that repeatedly arose in our listening sessions was the potential for AI to result in less respect for educators or less value for their skills. Across the nation, we are now responding to decreasing interest in entering or remaining in the teaching profession. Now is the time to show the respect and value we hold for educators by informing and involving them in every step of the process of designing, developing, testing, improving, adopting, and managing AI-enabled edtech. This includes involving educators in reviewing existing AI-enabled systems, tools, and data use in schools, designing new applications of AI based on teacher input, carrying out pilot evaluations of proposed new instructional tools, collaborating with developers to increase the trustworthiness of the deployed system, and raising issues about risks and unexpected consequences as the system is implemented.

We have already seen educators rise to the challenge of creating overall guidelines, designing specific uses of available AI-enabled systems and tools, and ferreting out concerns. And yet, the influence of educators in the future of AI-enabled products cannot be assumed; instead, constituents need policies that put muscle behind it. Could we create a national corps of leading educators representing every state and region to provide leadership? Could we commit to developing necessary professional development supports? Can we find ways to compensate educators so they can be at the forefront of designing the future of education? Our policies should enable educators to be closely involved in design of AI-enabled educational systems.

Although we know that the responsibility for informing and involving educators must be distributed at all levels of national and school governance, the Office of Educational Technology

can play a key role in informing and involving educators through its reports, events, outreach, and in a future NETP. Although examples above refer to K-12 teachers, higher education instructors must also be included. We also call on the edtech industry to involve educators throughout their design and development processes. For example, AI-enabled teaching assistants are only likely to help teachers do their job if teachers are thoroughly involved as the assistants are designed. We call upon institutions that prepare teachers to integrate technology more systematically into their programs; for example, the use of technology in teaching and learning should be a core theme across teacher preparation programs, not an issue that arises only in one course.

Recommendation #6: Focus R&D on Addressing Context and Enhancing Trust and Safety

Research that focuses on how AI-enabled systems can adapt to context (including variability among learners) in instructional approaches and across educational settings is essential to answering the question of, “Do specific applications of AI work in education, and if so, *for whom and under what conditions?*” The italicized phrase points to variability among learners and diversity in the settings for learning. We call upon innovators in R&D to focus their efforts to advance AI on the long tail of learning variability, where large populations of students would benefit from customization of learning. We also call on R&D to lead by establishing how trust can be strengthened in AI-enabled systems, building on the *Blueprint’s* call for safe and effective systems yet also including education-specific requirements, such as how teachers can be meaningfully involved in design phases, not only in implementation and evaluation.

Although many products today are adaptive, some adapt on just one or a few dimensions of variability, such as student’s accuracy in problem solving. As teachers know, there are many more important ways to adapt to students’ strengths and needs. Students are neurodiverse and may have specific disabilities. They bring different assets from their experiences at home, in communities, and in their cultures. They have different interests and motivations. They are in different places in their journeys to master the English language. And they learn in varied settings. Classrooms and schools are different, and at home, students learn in informal settings in ways that could complement school learning. We recommend attention to “context” as a means for expressing the multiple dimensions that must be considered when elaborating the phrase “for whom and under what conditions.” We also acknowledge the role of researchers in conducting evaluations, which must now consider not only efficacy but must also explore where harm may arise and the system problems that can occur through weak trust or over-trust in AI systems.

R&D must take the lead in making AI models more context-sensitive and ensuring that they are effective, safe, and trustworthy for use with varied learners in diverse settings. Although AI has capabilities to find patterns beyond the limited number of variables that people normally think about, AI is not particularly good at understanding and working with context in the ways people do. Over time, we’ve seen learning sciences grow to be less about individualistic cognitive principles and more encompassing first of social learning and then of the many dimensions of context that matter in learning. Our use of AI needs to follow this trajectory toward context to support educational applications.

To achieve human-centric vision, listening session attendees argued that teams will need time and freedom to explore how best to manage the tension between the pace of technological

advancement and the need for broader contextual insights—for trust and for safety. They will need time and freedom to pioneer new processes that better involve teachers and students as co-designers, with attention to balancing power dynamics. And they will need to shift attention from older ways of framing priorities (such as achievement gaps) to new ways of prioritizing digital equity. We call on R&D funders to focus resources on the long tail of learner variability, the need for AI-enabled systems that better incorporate context, and time required to get contextual considerations right. We call upon researchers and developers to prioritize challenges of context, trust, and safety in their work to advance AI.

Recommendation #7: Develop Education-Specific Guidelines and Guardrails

Our final recommendation is central to policymakers. A feature of the American educational system is the emphasis on local decision making. With technology growing in complexity at such a rapid pace, it is becoming difficult for local leaders to make informed decisions about the deployment of artificial intelligence. As we have discussed, the issues are not only data privacy and security but extend to new topics such as bias, transparency, and accountability. It will be harder to evaluate promising edtech platforms that rely on AI systems against this evolving, complex set of criteria.

Regulations related to key student and family data privacy laws like the Family Educational Rights & Privacy Act ([FERPA](#)), the Children's Internet Privacy Act ([CIPA](#)), and the Children's Online Privacy Protection Act ([COPPA](#)) warrant review and further consideration in light of new and emerging technologies in schools. Laws such as the Individuals with Disabilities Education Act ([IDEA](#)) may likewise be considered as new situations arise in the use of AI-enabled learning technologies. As discussed throughout this document, the *Blueprint for an AI Bill of Rights* is an important framework throughout this work.

The Department encourages parallel work by constituents in all levels of the educational system. In addition to the key federal laws cited immediately above, many states have also passed privacy laws that govern the use of educational technology and edtech platforms in classrooms. Further constituents can expect general frameworks for responsible AI in parallel sectors like health, safety, and consumer products to be informative but not sufficient for education's specific needs. Leaders at every level need awareness of how this work reaches beyond implications for privacy and security (e.g., to include awareness of potential bias and unfairness), and they need preparation to effectively confront the next level of issues.

Next Steps

We are heartened to see intensifying discussions throughout the educational ecosystem about the role of AI. We see progress that we can build upon occurring, as constituents discuss these three types of questions: What are the most significant opportunities and risks? How can we achieve trustworthy educational AI? How can we understand the models at the heart of applications of AI and ensure they have the qualities that align to educational aspirations?

The Department developed this report with awareness of contributions arising from many types of organizations and collectives. Internationally, we recognize parallel efforts to consider AI in the European Union, at the United Nations, and indeed throughout the world. We are aware of progress being led by organizations such as UNESCO, the EdSAFE AI Alliance, and research

organizations in many countries. We plan to continue cross-agency work, for example, by continuing to coordinate with the Office of Science and Technology Policy and other Federal agencies as agencies implement next steps guided by the *Blueprint for an AI Bill of Rights*. We see a broad and fertile context for necessary next steps:

- Working within this context and with others, the Department will consider specific policies and regulations so that educators can realize the opportunities of AI in edtech while minimizing risks. For example, the Department is developing a set of AI usage scenarios to strengthen the process of evaluating and enhancing policies and regulations. The principles and practices in the *Blueprint for an AI Bill of Rights* will be used to ensure the scenarios mitigate important risks and harms.
- Working with constituents (including education leaders; teachers, faculty, support staff, and other educators; researchers; policymakers; funders; technology developers; community members and organizations; and above all, learners and their families/caregivers), we will develop additional resources and events to increase understanding of AI and to involve those who will be most affected by these new technologies.
- Working across sectors, such as education, innovation, research, and policy, we will revise and update the NETP to guide all constituents toward safe, equitable, and effective AI in education in the United States, in alignment with our overall educational priorities.

Common Acronyms and Abbreviations

- AES: Automated Essay Scoring
- AI: Artificial Intelligence
- [CIPA](#): Children's Internet Protection Act
- [COPPA](#): Children's Online Privacy Protection Act
- Edtech: Educational Technology
- ESEA: Elementary and Secondary Education Act
- [ESSA](#): Every Student Succeeds Act
- [FERPA](#): [Family Educational Rights and Privacy Act](#)
- IA: Intelligence Augmentation
- IDEA: Individuals with Disabilities Education Act
- IEP: Individualized Education Program
- ITS: Intelligent Tutoring Systems
- [NETP](#): National Education Technology Plan
- R&D: Research & Development

Acknowledgements

Project Team

Artificial Intelligence and the Future of Teaching and Learning was developed under the leadership and guidance of **Roberto J. Rodríguez**, Assistant Secretary for the Office of Planning, Evaluation, and Policy Development, **Kristina Ishmael**, Deputy Director of the Office of Educational Technology, and **Bernadette Adams**, Senior Policy Advisor for the Office of Educational Technology at the U.S. Department of Education.

Support for the creation of this document was provided by Digital Promise, led by **Jeremy Roschelle** with **Carly Chillmon**, **Judi Fusco**, **Gabrielle Lue**, **Eric Nentrup**, **My Nguyen**, **Pati Ruiz**, and **Zohal Shah**. Special thanks to Center for Integrative Research in Computing and Learning Sciences postdocs **Michael Chang** and **Aditi Mallavarapu**.

Listening Session Panelists and Hosts

Hal Abelson	Judi Fusco	Aditi Mallavarapu
Ryan Baker	Dragan Gasevic	Ole Molvig
Nancye Blair Black	Kip Glazer	Peter Norvig
Marcelo Aaron Bonilla Worsley	Janice Gobert	Thomas Philip
Michael Chang	Sarah Hampton	Vidula Plante
Carly Chillmon	Kristina Ishmael	Jeremy Roschelle
Sherice Clarke	Jim Larimore	Pati Ruiz
Tammy Clegg	Nicol Turner Lee	Alina Von Davier
Sidney d'Mello	Sherry Loftin	Erin Walker
	Gabrielle Lue	Diego Zapata

We also thank 1,075 people who registered for Listening Sessions and 700 who attended.

References

- Akgun, S., Greenhow, C. (2022). Artificial intelligence in education: Addressing ethical challenges in K-12 settings. *AI Ethics*, 2, 431–440. <https://doi.org/10.1007/s43681-021-00096-7>
- Aleven, V., McLaughlin, E. A., Glenn, R. A., & Koedinger, K. R. (2016). Instruction based on adaptive learning technologies. In Mayer, R.E. & Alexander, P.A., *Handbook of research on learning and instruction*, 522-560. ISBN: 113883176X
- Baker, R.S., Esbenshade, L., Vitale, J., & Karumbaiah, S. (2022). Using demographic data as predictor variables: A questionable choice. <https://doi.org/10.35542/osf.io/y4wvj>
- Black, P. & Wiliam, D. (1998). Inside the black box: Raising standards through classroom assessment. *Phi Delta Kappan*, 92(1), 81-90. <https://kappanonline.org/inside-the-black-box-raising-standards-through-classroom-assessment/>
- Black, P., & Wiliam, D. (2009). Developing the theory of formative assessment. *Educational Assessment, Evaluation and Accountability*, 21(1), 5-31. <https://doi.org/10.1007/s11092-008-9068-5>
- Boden, M.A. (2018). *Artificial intelligence: A very short introduction*. Oxford. ISBN: 978-0199602919
- Bryant, J., Heitz, C., Sanghvi, S., & Wagle, D. (2020, January 14). *How artificial intelligence will impact K-12 teachers*. McKinsey. <https://www.mckinsey.com/industries/education/our-insights/how-artificial-intelligence-will-impact-k-12-teachers>
- Celik, I., Dindar, M., Muukkonen, H. & Järvelä, S. (2022). The promises and challenges of artificial intelligence for teachers: A systematic review of research. *TechTrends*, 66, 616–630. <https://doi.org/10.1007/s11528-022-00715-y>
- Center for Integrative Research in Computing and Learning Sciences (CIRCLS). (2022, Feb.). *From Broadening to empowering: Reflecting on the CIRCLS'21 Convening*. <https://circls.org/circls21report>
- Chen, C., Park, H.W. & Breazeal, C. (2020). Teaching and learning with children: Impact of reciprocal peer learning with a social robot on children's learning and emotive engagement. *Computers & Education*, 150, <https://doi.org/10.1016/j.compedu.2020.103836>
- Chen, G., Clarke, S., & Resnick, L.B. (2015). Classroom Discourse Analyzer (CDA): A discourse analytic tool for teachers. *Technology, Instruction, Cognition and Learning*, 10(2), 85-105
- Dieterle, E., Dede, C. & Walker, M. (2022). The cyclical ethical effects of using artificial intelligence in education. *AI & Society*. <https://link.springer.com/article/10.1007/s00146-022-01497-w>
- Doewes, A. & Pechenizkiy, M. (2021). On the limitations of human-computer agreement in automated essay scoring. In *Proceedings of the 14th International Conference on Educational Data Mining (EDM21)*. https://educationaldatamining.org/EDM2021/virtual/static/pdf/EDM21_paper_243.pdf
- Englebart, D.C. (October 1962). *Augmenting human intellect: A conceptual framework*. SRI Summary Report AFOSR-3223. <https://www.dougenelbart.org/pubs/augment-3906.html>
- Ersozlu, Z., Ledger, S., Ersozlu, A., Mayne, F., & Wildy, H. (2021). Mixed-reality learning environments in teacher education: An analysis of TeachLivETM Research. *SAGE Open*, 11(3). <https://doi.org/10.1177/21582440211032155>
- European Commission, Directorate-General for Education, Youth, Sport and Culture. (2022). *Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators*. Publications Office of the European Union. <https://data.europa.eu/doi/10.2766/153756>
- Forsyth, S., Dalton, B., Foster, E.H., Walsh, B., Smilack, J., & Yeh, T. (2021, May). Imagine a more ethical AI: Using stories to develop teens' awareness and understanding of artificial intelligence and its societal impacts. In *2021 Conference on Research in Equitable and Sustained Participation in Engineering, Computing, and Technology (RESPECT)*. IEEE. <https://doi.org/10.1109/RESPECT51740.2021.9620549>
- Friedman, L., Blair Black, N., Walker, E., & Roschelle, J. (November 8, 2021) Safe AI in education needs you. *Association of Computing Machinery BLOG@ACM*, <https://cacm.acm.org/blogs/blog-cacm/256657-safe-ai-in-education-needs-you/fulltext>

- Gardner, J., O'Leary, M. & Yuan, L. (2021). Artificial intelligence in educational assessment: "Breakthrough? Or buncombe and ballyhoo?" *Journal of Computer Assisted Learning*, 37(5), 1207–1216. <https://doi.org/10.1111/jcal.12577>
- Gartner (n.d.) *Gartner glossary: Augmented intelligence*. Gartner. <https://www.gartner.com/en/information-technology/glossary/augmented-intelligence>
- Gay, G. (2018). *Culturally responsive teaching: Theory, research, and practice*. Teachers College Press. ISBN: 978-0807758762
- Godwin-Jones, R. (2021). Big data and language learning: Opportunities and challenges. *Language Learning & Technology*, 25(1), 4–19. <http://hdl.handle.net/10125/44747>
- Hammerness, K., Darling-Hammond, L., & Bransford, J. (2005). *Preparing teachers for a changing world: What teachers should learn and be able to do*. Jossey-Bass. ISBN: 0787996343
- Holmes, W. & Porayska-Pomsta, K. (Eds.) (2022). *The ethics of artificial intelligence in education*. Routledge. ISBN 978-0367349721
- Holstein, K., McLaren, B.M., & Aleven, V. (2019). Co-designing a real-time classroom orchestration tool to support teacher–AI complementarity. *Journal of Learning Analytics*, 6(2). <https://doi.org/10.18608/jla.2019.62.3>
- IEEE-USA Board of Directors. (February 10, 2017). *Artificial intelligence research, development and regulation*. IEEE <http://globalpolicy.ieee.org/wp-content/uploads/2017/10/IEEE17003.pdf>
- Jensen, E., Dale, M., Donnelly, P.J., Stone, C., Kelly, S., Godley, A. & D'Mello, S.K. (2020). Toward automated feedback on teacher discourse to enhance teacher learning. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (CHI '20). <https://doi.org/10.1145/3313831.3376418>
- Kai, S., Almeda, M.V., Baker, R. S., Heffernan, C., & Heffernan, N. (2018). Decision tree modeling of wheel-spinning and productive persistence in skill builders. *Journal of Educational Data Mining*, 10(1), 36–71. <https://doi.org/10.5281/zenodo.3344810>
- Kaplan, R.M., & Saccuzzo, D.P. (2017). *Psychological testing: Principles, applications, and issues*. Cengage Learning.
- Ke, Z., & Ng, V. (2019). Automated essay scoring: A survey of the state of the art. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 6300–6308. <https://doi.org/10.24963/ijcai.2019/879>
- Khosravi, H., Shum, S.B., Chen, G., Conati, C., Tsai, Y.-S., Kay, J., Knight, S., Martinez-Maldonado, R., Sadiq, S., Gašević, D. (2022). Explainable artificial intelligence in education. *Computers and Education: Artificial Intelligence*, 3. <https://doi.org/10.1016/j.caeai.2022.100074>
- Kulik, J.A., & Fletcher, J.D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*, 86(1), 42–78
- Ma, W., Adescope, O.O, Nesbit, J.C. & Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology*, 106(4), 901–918. <http://dx.doi.org/10.1037/a0037123>
- Maslej, N., Fattorini, L., Brynjolfsson E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Ngo, H., Niebles, J.C., Parli, V., Shoham, Y., Wald, R., Clark, J. and Perrault, R., (2023). *The AI index 2023 annual report*. Stanford University: AI Index Steering Committee, Institute for Human-Centered AI.
- Merrill, S. (2020). *In schools, are we measuring what matters?* Edutopia. <https://www.edutopia.org/article/schools-are-we-measuring-what-matters>
- Molenaar, I. (2022). Towards hybrid human-AI learning technologies. *European Journal of Education*, 00, 1–14. <https://doi.org/10.1111/ejed.12527>
- Mostow, J., Aist, G., Burkhead, P., Corbett, A., Cuneo, A., Eitelman, S., Huang, C., Junker, B., Sklar, M.B., & Tobin, B. (2003). Evaluation of an automated reading tutor that listens: Comparison to human tutoring and classroom instruction. *Journal of Educational Computing Research*, 29(1), 61–117. <https://doi.org/10.2190/06AX-QW99-EQ5G-RDCF>
- Mousavinasab, E., Zarifsanaiy, N., R. Niakan Kalhori, S., Rakhshan, M., Keikha, L., & Ghazi Saeedi, M. (2021). Intelligent tutoring systems: A systematic review of characteristics, applications, and evaluation methods. *Interactive Learning Environments*, 29(1), 142–163. <https://psycnet.apa.org/doi/10.1080/10494820.2018.1558257>
- National Academies of Sciences, Engineering, and Medicine. 2018. *How people learn II: Learners, contexts, and cultures*. The National Academies Press. <https://doi.org/10.17226/24783>

- National Research Council. 2000. *How people learn: Brain, mind, experience, and school*. The National Academies Press. <https://doi.org/10.17226/9853>
- Nentrup, E. (2022). *How Policymakers Can Support Educators and Technology Vendors Towards SAFE AI*. EdSAFE AI Alliance. <https://www.edsafeai.org/post/how-policymakers-can-support-ai>
- Page, E.B. (1966). The imminence of grading essays by computer. *Phi Delta Kappan*, 47(5), 238–243
- Paris, D., & Alim, H.S. (Eds.). (2017). *Culturally sustaining pedagogies: Teaching and learning for justice in a changing world*. Teachers College Press. ISBN: 978-0807758342
- Plass, J.L., & Pawar, S. (2020). Toward a taxonomy of adaptivity for learning. *Journal of Research on Technology in Education*, 52(3), 275–300. <https://doi.org/10.1080/15391523.2020.1719943>
- Regona, Massimo & Yigitcanlar, Tan & Xia, Bo & Li, R.Y.M. (2022). Opportunities and adoption challenges of AI in the construction industry: A PRISMA review. *Journal of Open Innovation Technology Market and Complexity*, 8(45). <https://doi.org/10.3390/joitmc8010045>
- Reynolds, C.R., & Suzuki, L.A. (2012). Bias in psychological assessment: An empirical review and recommendations. *Handbook of Psychology, Second Edition*. <https://doi.org/10.1002/9781118133880.hop210004>
- Ritter, S., Anderson, J.R., Koedinger, K.R. & Corbett, A. (2007). Cognitive Tutor: Applied research in mathematics education. *Psychonomic Bulletin & Review*, 14, 249–255/ <https://doi.org/10.3758/BF03194060>
- Roll, I., Aleven, V., McLaren, B.M., Koedinger, K.R. (2011). Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system, *Learning and Instruction*, 21(2), 267–280. <https://doi.org/10.1016/j.learninstruc.2010.07.004>.
- Roschelle, J., Dimitriadis, Y. & Hoppe, U. (2013). Classroom orchestration: Synthesis. *Computers & Education*, 69, 512–526. <https://doi.org/10.1016/j.compedu.2013.04.010>
- Roschelle, J., Feng, M., Murphy, R. & Mason, C.A. (2016). Online mathematics homework increases student achievement. *AERA Open*, 2(4), 1–12. DOI: 10.1177/2332858416673968
- Roschelle, J., Penuel, W., & Shechtman, N. (2006). Co-design of innovations with teachers: definition and dynamics. In *Proceedings of the 7th International Conference on Learning Sciences*, Bloomington, IN. <https://doi.org/10.22318/icls2006.606>
- Rose, D. (2000). Universal design for learning. *Journal of Special Education Technology*, 15(4), 47–51. <https://doi.org/10.1177/016264340001500407>
- Ruiz, P. & Fusco, J. (2022). *Teachers partnering with artificial intelligence: Augmentation and automation*. Digital Promise. <https://digitalpromise.org/2022/07/06/teachers-partnering-with-artificial-intelligence-augmentation-and-automation/>
- Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking. ISBN 978-0-525-55861-3.
- Shao, Q., Sniffen, A., Blanchet, J., Hillis, M.E., Shi, X., Haris, T.K., & Balkcom, D. (2020). Teaching american sign language in mixed reality. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(4), 1–27. <https://doi.org/10.1145/3432211>
- Sharples, M. & Pérez y Pérez, R. (2022). *Story machines: How computers have become creative writers*. Routledge. ISBN 9780367751951
- Shemshack, A., Spector, J.M. (2020) A systematic literature review of personalized learning terms. *Smart Learning Environments*, 7(33). <https://doi.org/10.1186/s40561-020-00140-9>
- Shute, V.J. (2008). Focus on formative feedback. *Review of Educational Research*, 78(1), 153–189. <https://doi.org/10.3102/0034654307313795>
- Shute, V. J., Ventura, M., & Kim, Y. J. (2013). Assessment and learning of qualitative physics in Newton's Playground. *The Journal of Educational Research*, 106(6), 423–430. <https://doi.org/10.1080/00220671.2013.832970>
- Swiecki, Z., Khosravi, H., Chen, G., Martinez-Maldonado, R., Lodge, J.M., Milligan, S., Selwyn, B. & Gašević, D. (2022). Assessment in the age of artificial intelligence. *Computers and Education: Artificial Intelligence*, 3. <https://doi.org/10.1016/j.caeai.2022.100075>
- The White House (February 17, 2023). *Executive order on further advancing racial equity and support for underserved communities through the federal government*. <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/02/16/executive-order-on-further-advancing-racial-equity>

- The White House (September 8, 2022). *Readout of White House listening session on tech platform accountability*. <https://www.whitehouse.gov/briefing-room/statements-releases/2022/09/08/readout-of-white-house-listening-session-on-tech-platform-accountability/>
- U.S. Department of Education, Office of Educational Technology (2022). *Advancing digital equity for all: Community-based recommendations for developing effective digital equity plans to close the digital divide and enable technology-empowered learning*. US Department of Education.
- U.S. Department of Education, Office of Educational Technology. (2010). *Transforming American Education: Learning Powered by Technology*. U.S. Department of Education. p. 78
- Van Lehn, K. (2011) The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197-221. <https://doi.org/10.1080/00461520.2011.611369>
- Wagner, A.R., Borenstein, J. & Howard, A. (September 2018). Overtrust in the robotics age. *Communications of the ACM*, 61(9), 22-24. <https://doi.org/10.1145/3241365>
- Walker, E., Rummel, N. & Koedinger, K.R. (2015). Adaptive intelligent support to improve peer tutoring in algebra. *International Journal of Artificial Intelligence in Education*, 24, 33-61 <https://doi.org/10.1007/s40593-013-0001-9>
- Walton Family Foundation (March 1, 2023). *Teachers and students embrace ChatGPT for education*. <https://www.waltonfamilyfoundation.org/learning/teachers-and-students-embrace-chatgpt-for-education>
- Webb, N.M., & Farivar, S. (1994). Promoting helping behavior in cooperative small groups in middle school mathematics. *American Educational Research Journal*, 31(2), 369-395. <https://doi.org/10.3102/00028312031002369>
- White House Office of Science and Technology Policy (October 2022), *Blueprint for an AI bill of rights: Making automated systems work for the American people*. The White House Office of Science and Technology Policy. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
- Wiggins, G. (2015). *Seven keys to effective feedback*. ASCD. <https://www.ascd.org/el/articles/seven-keys-to-effective-feedback>
- Winne, P.H. (2021). Open learner models working in symbiosis with self-regulating learners: A research agenda. *International Journal of Artificial Intelligence in Education*, 31(3), 446-459. <https://doi.org/10.1007/s40593-020-00212-4>
- Zacamy, J. & Roschelle, J. (2022). *Navigating the tensions: How could equity-relevant research also be agile, open, and scalable?* Digital Promise. <http://hdl.handle.net/20.500.12265/159>
- Zhai, X., He, P., Krajcik, J. (2022). Applying machine learning to automatically assess scientific models. *Journal of Research in Science Teaching*. <https://doi.org/10.1002/tea.21773>
- Zhang, H., Lee, I., Ali, S., DiPaola, D., Cheng, Y., & Breazeal, C. (2022). Integrating ethics and career futures with technical learning to promote AI literacy for middle school students: An exploratory study. *International Journal of Artificial Intelligence in Education*, 1-35. <https://doi.org/10.1007/s40593-022-00293-3>