Artificial Intelligence in Education

Tatenda Whitney Mavhunga[1, a]

[1] Xi'an Technological University, Xi'an, Shaanxi, China

[a]whitneymav@gmail.com

Abstract. This research paper delves into the incorporation of Artificial Intelligence (AI) in the educational sector, analyzing its advantages, obstacles, and future possibilities. AI holds the promise of transforming education through personalized learning, enhanced teaching techniques, and streamlined administrative tasks. Nonetheless, its adoption also brings forth ethical, privacy, and equity issues. The paper surveys current AI applications in education, assesses their effectiveness, and explores the implications for educators and policymakers.

Keywords: AI (artificial intelligence); Family Educational Rights and Privacy Act (FERPA); Intelligence Augmentation (IA); Intelligent Tutoring System (ITS); Individualized Education Programs (IEPs).

1. Introduction

AI can be described as "automation based on associations." When computers automate reasoning based on associations found in data (or associations derived from expert knowledge), two fundamental shifts in AI occur, advancing computing beyond conventional edtech: transitioning from merely capturing data to detecting patterns within it, and moving from providing access to instructional resources to automating decisions about instruction and other educational processes. These capabilities of detecting patterns and automating decisions represent significant advancements in the responsibilities that can be assigned to a computer system. However, the development of AI systems can introduce biases in pattern detection and unfairness in automated decision-making. Therefore, it is crucial for educational systems to regulate their use of AI.

1.1 Rising Interest in AI in Education. Today, many priorities for improving teaching and learning remain unmet. Educators are seeking technology-enhanced approaches that are safe, effective, and scalable to address these needs. Naturally, they are curious whether the rapid advances in everyday technology can help. Like everyone else, educators use AI-powered services in their daily lives, such as home voice assistants, grammar and sentence completion tools, and automated trip planning apps. They see potential in using AI capabilities like speech recognition to support students with disabilities, multilingual learners, and others who could benefit from more adaptive and personalized digital learning tools.

However, educators are also aware of new risks. The powerful functionality of AI comes with data privacy and security concerns. They recognize that AI can sometimes generate inappropriate or incorrect output and have noted new ways students might misrepresent others' work as their own.

In late 2022 and early 2023, the public became aware of new generative AI chatbots and began exploring their potential for writing essays, creating lesson plans, producing images, and generating personalized student assignments.

1.2 Three Reasons to Address AI in Education Now. First, AI has the potential to meet educational priorities more effectively, at scale, and at lower costs. Addressing the varied unfinished learning of students due to the pandemic is a key policy priority, and AI can enhance the adaptivity of learning resources to meet students' strengths and needs. Improving teaching jobs is also a priority; AI, through automated assistants and other tools, can provide teachers with greater support and enable them to extend individual support to students when time is limited.

Second, there is an urgency and importance driven by awareness of system-level risks and anxiety about potential future risks. For example, increased surveillance of students is a concern. Some teachers worry about being replaced by AI; however, the Department firmly rejects this notion. Public concerns about algorithmic bias, such as voice recognition systems struggling with

regional dialects or exam monitoring systems unfairly identifying certain groups of students for disciplinary action, highlight the potential for discrimination.

Third, the urgency is compounded by the scale of possible unintended or unexpected consequences. When AI automates instructional decisions at scale, it may lead to unwanted outcomes. For instance, if AI adjusts the curriculum pace differently for students based on incomplete data, flawed theories, or biased assumptions, achievement gaps could widen.

1.3 Toward Policies for AI in Education. In the United States, several executive orders aim to ensure AI is trustworthy and fair, supported by the White House Office of Science and Technology Policy's Blueprint, which outlines relevant principles and practices. These initiatives, along with other AI-related policy efforts in both the executive and legislative branches, will guide AI use across all sectors of society. In Europe, the European Commission recently issued ethical guidelines on using AI and data in education for teachers.

Policies must consider existing frameworks, including federal student privacy laws like the Family Educational Rights and Privacy Act (FERPA) and similar state laws.

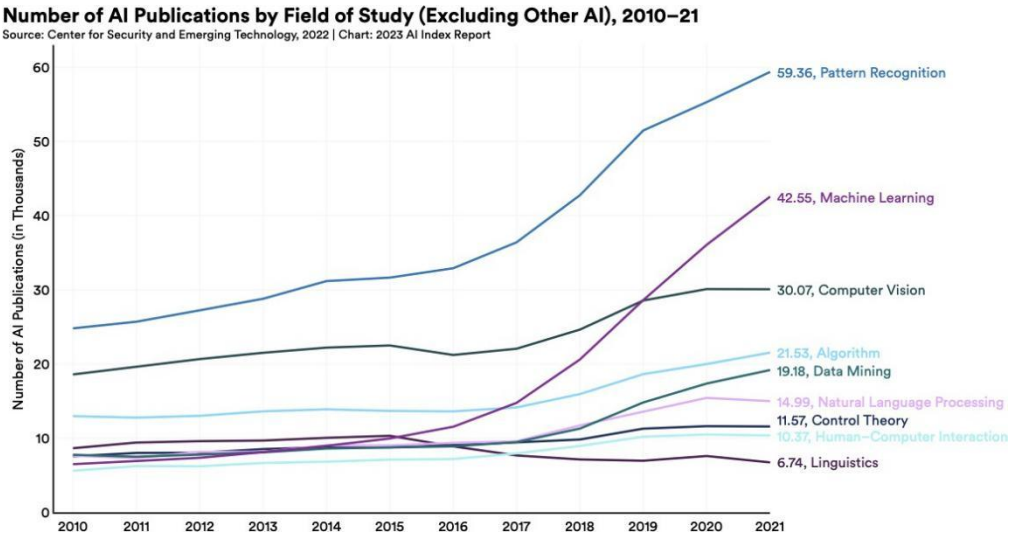

Figure 1. Research about AI is growing rapidly. Other indicators, such as dollars invested and number of people employed, show similar trends

AI is advancing rapidly (see Figure 1), bringing powerful new features for generating images and text to the public and transforming how people create content. These advancements are not confined to research labs but are also making headlines in mainstream media and educational publications.

There is an urgent need for policies to:

1. Utilize automation to enhance learning outcomes while preserving human decision-making and judgment.

2. Examine the quality of data in AI models to ensure fair and unbiased pattern recognition and decision-making in educational applications, based on accurate and context-appropriate information.

2. Building Ethical, Equitable Policies Together

2.1 Guiding Questions. Recognizing that AI enhances automation and enables machines to perform tasks previously done solely by humans brings us to two fundamental questions:

1. What is our shared vision for an educational system that effectively uses automation to improve learning while safeguarding and prioritizing human agency?

2. How and when will we establish the necessary guidelines and safeguards, along with compelling evidence of positive impacts, to ensure this vision can be ethically and equitably implemented on a broad scale?

2.2 Foundation 1: Center People (Parents, Educators, and Students). AI-focused education policies at the federal, state, and district levels are essential to guide and empower local and individual decisions regarding technology adoption and use in schools and classrooms. Many people rely on AI-enabled products for their effectiveness and convenience; for instance, few prefer paper maps because technology provides more efficient route planning. However, the privacy trade-offs involved in using AI systems are often overlooked.

2.3 Foundation 2: Advance Equity. This report defines digital equity as a state where individuals and communities have the necessary information technology resources to fully participate in society and the economy. We noticed an increasing focus on data quality issues and the consequences of using poor or inappropriate data in AI systems for education. AI development depends on datasets, and if these datasets are non-representative or contain undesirable associations, the resulting AI models may unfairly detect patterns or make automated decisions. This systematic unfairness, known as "algorithmic bias," can undermine equity by causing unintended discrimination.

Algorithmic bias is intrinsic to the development of AI algorithms using historical data, making it difficult to anticipate all the impacts of biased data and algorithms during system design. In postsecondary education, for example, algorithms that make enrollment decisions, identify students for early intervention, or flag potential cheating must be examined for unfair discriminatory bias—not only during the initial design but also as these systems are implemented and widely used.

2.4 Foundation 3: Ensure Safety, Ethics, and Effectiveness. The development and implementation of AI necessitates access to comprehensive data. This data extends beyond traditional student records, such as rosters and grade books, to include detailed information about students' interactions with technology during learning and teachers' use of technology in their teaching practices.

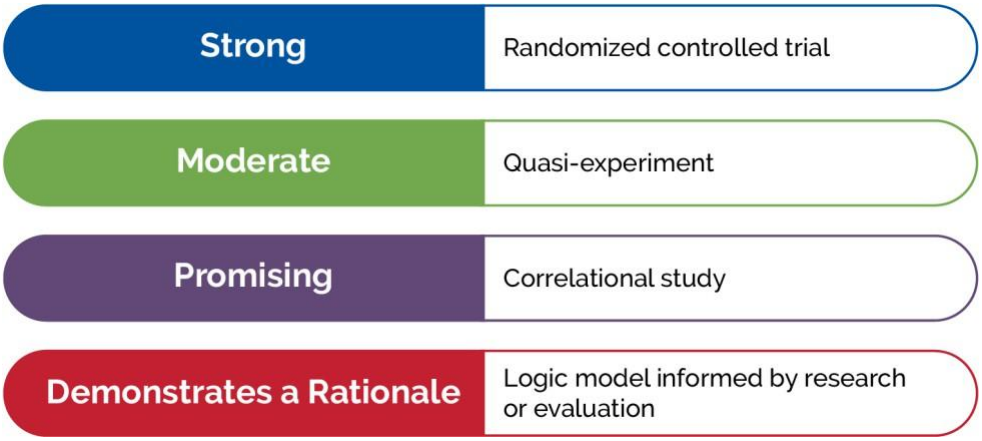

Figure 2. The Elementary and Secondary Education Act defines four levels of evidence

2.5 Foundation 4: Promote Transparency. In education, decision makers require more than just notifications; they need a deep understanding of how AI models function across various educational scenarios.

Since AI models in edtech are approximations of reality, stakeholders should always consider these questions: How precise are the AI models? What are the broader implications of using AI models at scale in educational processes?

3 What is AI?

AI is not a single but an umbrella term encompassing an expanding range of modeling capabilities, as visualized in Figure 3.

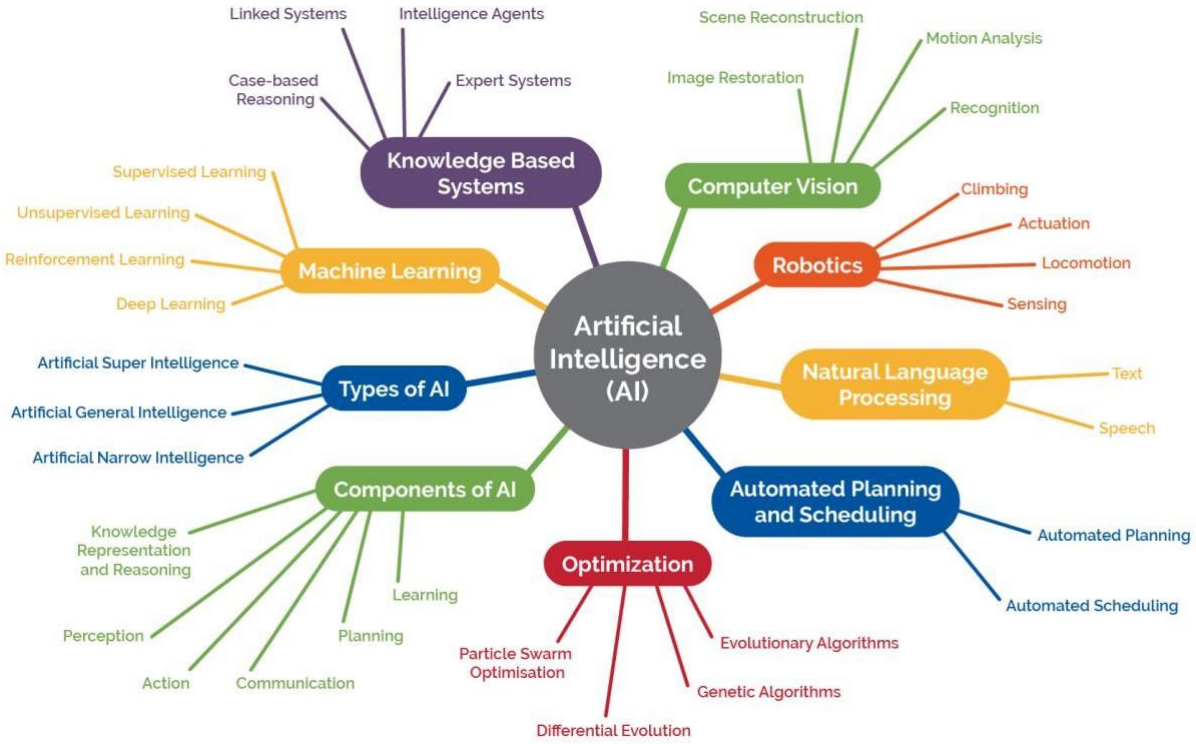

Figure 3. Components, types, and subfields of AI based on Regona et al (2022)

3.1 Perspective: Human-Like Reasoning. The concept of "human-like" is useful as shorthand for indicating that modern computers possess capabilities far beyond those of early edtech applications. Educational tools will be able to engage in conversations with students and teachers, assist in managing classroom activities, and take actions that have broader impacts on both students and teachers.

3.2 Perspective: An Algorithm that Pursues a Goal. This second definition highlights that AI systems and tools identify patterns and make decisions to achieve specific goals. These pattern recognition capabilities and automated recommendations will influence the educational process, impacting both student learning and teacher decision-making. For example, today's personalized learning systems can detect when a student is struggling and suggest an alternative instructional sequence. The range of pattern recognition and automated recommendations will continue to grow. For decades, the boundaries between the roles of teachers and computers have been discussed in education, with terms like "computer-aided instruction," "blended instruction," and "personalized learning." But how are instructional choices made in systems involving both humans and algorithms? Currently, AI systems and tools enable the adaptation of instructional sequences to student needs, providing feedback and hints during activities like mathematics problem-solving or foreign language learning.

Consider another simple example. When a teacher requests, "Display a map of ancient Greece on the classroom screen," an AI system might select from hundreds of maps by considering the lesson objectives, successful outcomes in similar classrooms, or maps with features beneficial for student learning. In this scenario, an AI system's suggestion of instructional resources or options can save the instructor time, allowing them to focus on more important goals. However, there are instances where an instructor might reject AI-enabled automation, such as allowing an AI system to choose the most appropriate and relevant readings for students related to a historical event.

Computers process theory and data differently than humans. AI's success hinges on the associations or relationships identified in the data fed to an algorithm during the AI model development process. While some associations can be useful, others may be biased or inappropriate, posing a significant risk of algorithmic discrimination.

Humans bring a unique perspective to agency, goal-setting, and reasoning, utilizing their ability to understand multiple contexts. For instance, a teacher might notice that three students make the same mathematical error but realize that one has an Individualized Education Program for vision issues, another misunderstands a mathematical concept, and the third just had a frustrating interaction on the playground; thus, the same instructional decision wouldn't be suitable for all.

3.3 Perspective: Intelligence Augmentation. Intelligence Augmentation (IA) emphasizes human intelligence and decision-making while acknowledging that people can be overwhelmed and benefit from supportive tools. AI can assist teachers in making more informed decisions by identifying patterns that might go unnoticed by humans.

3.4 Definition of "Model". To meaningfully assess AI, stakeholders must consider specific models and their development processes. AI models, like financial models, approximate reality to identify patterns, make predictions, or analyze different decisions. For example, in a typical middle school math curriculum, students use a mathematical model to determine the better of two cell phone plans. Similarly, financial planners use such models to advice on retirement portfolios. Essentially, AI is a sophisticated mathematical toolkit for building and utilizing models.

3.5 Insight: AI Systems Enable New Forms of Interaction. AI models enable computational processes to make recommendations and plans, as well as support more natural forms of interaction, such as conversing with an assistant. Teachers and students can interact with AI by selecting items from menus, answering multiple-choice questions, typing short responses, dragging objects, or using touch gestures. The computer responds through text, graphics, and multimedia. While versatile, these interactions are distinctly human-computer and differ from human-to-human interaction. With AI advancements, computer interactions are becoming more human-like (see Figure 4). For example, a teacher might speak to an AI assistant that responds verbally, or a student might draw something that the computer enhances. A teacher or student might start writing, and the AI could complete their sentence, similar to how modern email programs predict and finish thoughts.

	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

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

4. Learning

Students engage in discussions that deepen their understanding, use visualizations and simulations to explain real-world concepts, and benefit from supportive scaffolding and timely feedback.

4.1 Insight: AI Enables Adaptivity in Learning. AI can enhance technology's ability to meet students where they are, build on their strengths, and help them grow their knowledge and skills.

This is possible due to AI's capability to work with natural forms of input and the robust foundation of AI models. However, the core limitations stem from the nature of the models at the heart of any AI-enabled system. These models are approximations of reality. If crucial aspects of human learning are left out or underdeveloped in the model, the adaptivity and support for learning will be limited and potentially fragile.

AI models have shown greater capabilities due to advances in "large language models" or "foundational models." These general models, despite their impressive skills, still have limitations. For instance, generative AI models can quickly create convincing essays on a variety of topics or generate credible images from a few prompts. Despite the excitement surrounding foundational models, they are still narrower than the broader vision for human learning, making it crucial to design learning environments with these limitations in mind.

4.2 Intelligent Tutoring Systems: An Example of AI Models. One enduring type of AI technology is the Intelligent Tutoring System (ITS). In an initial breakthrough, researchers successfully developed precise models of how proficient individuals solve mathematical problems. These models were integrated into a system that observed students as they tackled mathematical challenges on a computer. Studies examining human tutors have highlighted that providing feedback on specific steps, rather than merely indicating correctness, is likely a crucial factor in effective tutoring. For instance, when a student deviated from the expert model, the system offered guidance to help them realign with the correct approach.

Over time, these initial limitations have been addressed through two main approaches: enhancing AI models and incorporating human oversight—a perspective that remains crucial today. For instance, in contemporary ITS applications, while the system may specialize in delivering feedback during practice sessions, human teachers remain responsible for fostering student engagement, promoting self-regulation, and managing other instructional aspects.

4.3 Important Directions for Expanding AI-Based Adaptivity. The term "adaptivity" is often synonymous with "personalization." However, many observers have criticized its lack of precision. Some educators interpret personalization as empowering learners with "voice and choice," while others see it as a learning management system recommending a tailored "playlist" of activities for each student. For instance, skilled teachers adeptly engage students by drawing on their personal experiences and refine their explanations until a profound "aha!" moment occurs.

There are many ways in which the core models in an AI system must be expanded.

1. Shifting from deficit-based to asset-oriented. Historically, adaptive systems have identified and addressed student weaknesses. Moving forward, it's crucial for AI models to recognize and capitalize on students' strengths or competencies to enhance their learning experiences.

2. Expanding from individual cognition to including social and other aspects of learning dimensions. Current adaptive learning paradigms primarily emphasize personalized cognitive learning, often sidelining social and motivational aspects. A more comprehensive approach acknowledges the importance of social interaction and diverse learning contexts.

3. Embracing from neurodiversity learning models. AI systems should accommodate neurodiverse learners who may require alternative learning paths and interfaces tailored to their strengths. This inclusivity extends to learners with disabilities, promoting multiple learning routes and interactive modalities.

4. Transitioning from fixed tasks to dynamic, open-ended, and challenges. Traditional AI models excel at structured tasks like solving mathematical problems but often struggle with fostering creativity and prolonged engagement in open-ended tasks. To support lifelong learning and holistic development, AI systems should empower students to innovate and tackle creative challenges, encouraging active participation and innovative problem-solving skills.

5. Moving beyond correct answers to broader learning objectives. Many current adaptive systems prioritize correct responses, adjusting learning pace based on accuracy. However, effective learning goes beyond mere correctness, encompassing skills such as self-regulation, persistence in problem-solving, teamwork, and goal-setting autonomy.

4.4 A Challenge: Systems Thinking About AI in Education. As previously discussed, AI systems and tools do not fully align with educational goals, requiring us to design educational environments that appropriately integrate AI, allowing educators and other adults to use these tools effectively for teaching and learning. Additionally, small-group work is likely to remain important, as students might collaborate in small groups to use mathematics to predict or justify responses to realistic challenges.

5. Teaching

It's also crucial to remember the reasons people pursue a career in teaching and ensure they can focus on what truly matters.

This section explores how AI can support teachers and enhance education, including:

1. AI assistants that alleviate routine teaching tasks.
2. AI systems offering tailored recommendations to address students' needs and augment teachers' efforts.
3. AI tools that help teachers reflect on their practices, plan lessons, and improve their teaching strategies.

5.1 Always Center Educators in Instructional Loops. To successfully integrate AI as a tool for enhancing learning and teaching, it is essential to always prioritize educators (ACE).

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

Figure 5 highlights three critical loops, inspired by research on adaptivity loops:

1. The loop where teachers make real-time decisions during the act of teaching.
2. The loop where teachers engage in planning, preparation, and reflection on their teaching practices, including professional development.
3. The loop where teachers are involved in decisions regarding the design, selection, and evaluation of AI-enabled technologies, thereby influencing not only their own classrooms but also those of their peers.

5.2 Insight: Using AI to Improve Teaching Jobs. Teaching is an infamously complex profession, requiring teachers to make thousands of decisions daily. Additionally, they are integral members of their communities, expected to engage with families and caregivers.

Consider how much simpler some everyday tasks have become. We can now request and receive alerts and notifications about various events.

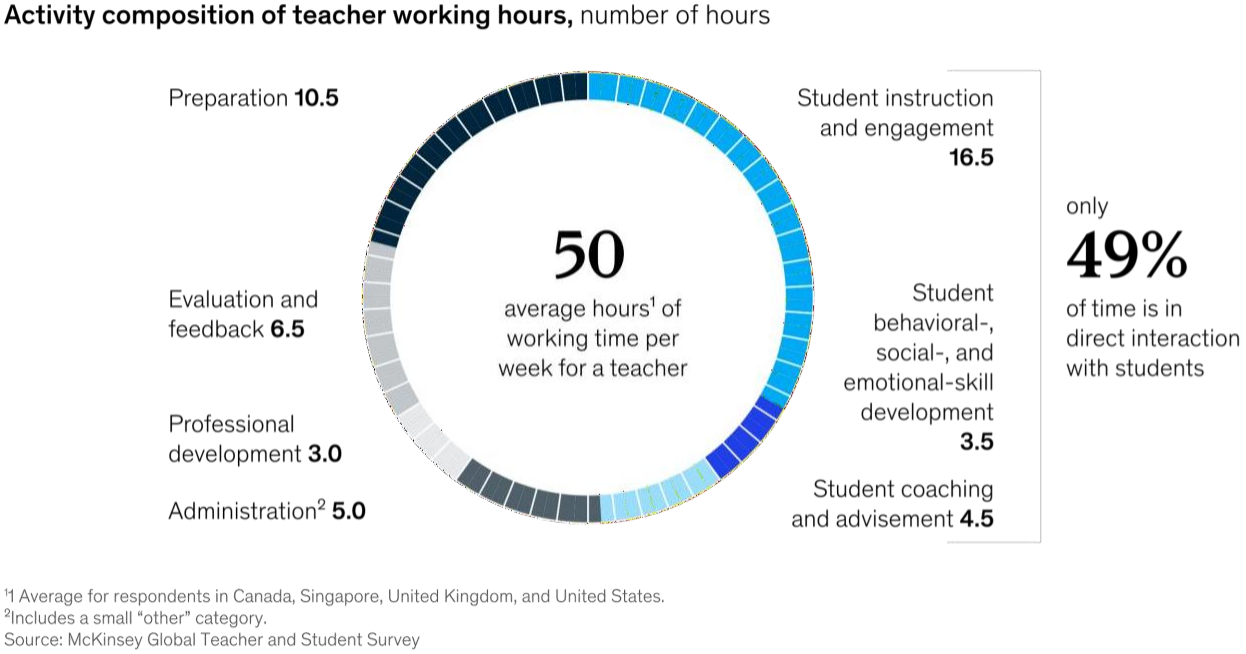

Figure 6. Teachers work about 50 hours a week, spending less than half the time in direct interaction with students.

The time saved through AI-enabled technology should be redirected towards more effective instruction. For example, reducing the average weekly preparation time from eleven hours to just six.

We highlight these opportunities, below.

1. Handling low-level details to ease teaching burdens and increase focus on students. A good teacher must master details both large and small. For instance, when working with a particular student, the teacher might want to send that student a helpful learning resource later. How will they remember to send it? A voice assistant or other AI assistant could help by organizing simple voice notes for teachers to follow up on after class. AI-enabled voice assistants are beginning to emerge in the market, capable of handling various simple tasks to help teachers stay focused on their students. These tasks include record-keeping, starting and stopping activities, controlling displays, speakers, and other classroom technologies, and providing reminders.

2. Extending beyond the teacher's availability with their students but continuing to deliver on the teacher's intent. Teachers often wish they could do more for each student but are limited by the number of hours in a school day. For instance, a teacher might want to sit with a student as they practice ten additional math problems, providing continuous support and feedback. If the teacher can only assist with three problems, they could delegate the rest to an AI-enabled learning system. Teachers can't be available around the clock to help with homework, but they can specify the types of support, hints, and feedback they want students to receive during after-school study sessions. An AI assistant can provide students with support anytime and anywhere they do homework or practice skills independently.

3. Making teacher professional development more productive and fruitful. AI can also compute metrics, such as whether students are talking more or less, which are difficult for a teacher to assess during a lesson.

For educators aiming to enhance student engagement, these metrics offer significant utility. Additionally, there is a growing availability of classroom simulation tools that allow teachers to refine their teaching abilities in authentic scenarios. These simulators incorporate real classroom teaching examples, with the ability to alter participant faces and voices, facilitating the sharing and discussion of teaching experiences among educators while maintaining anonymity.

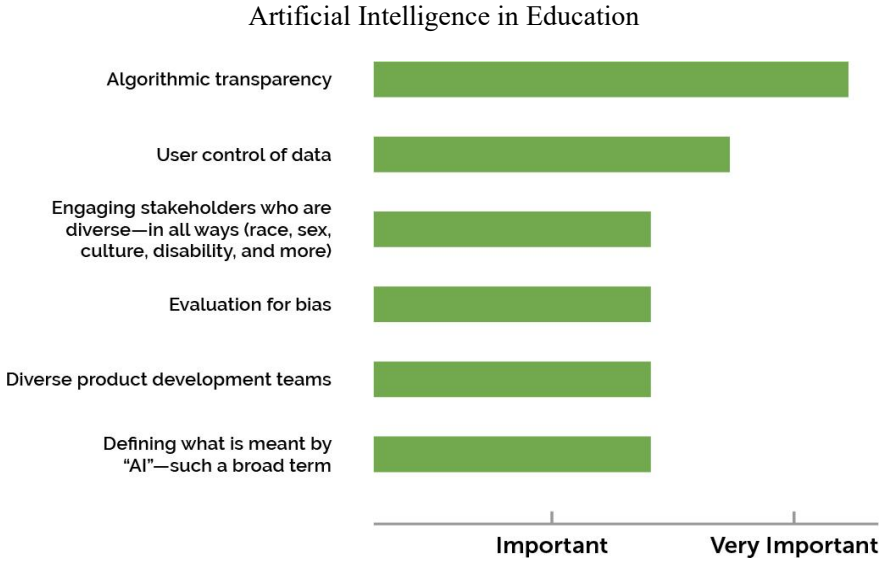

Figure 7. Concerns raised during the listening session about teaching with AI

5.3 Preparing and Supporting Teachers in Planning and Reflecting. In-service teachers will similarly require ongoing professional development to harness the potential benefits of AI.

Professional development efforts must strike a balance: not only exploring opportunities but also educating teachers about potential risks associated with AI, while equipping them with tools to navigate these challenges effectively.

AI has the potential to assist teachers by offering personalized recommendations tailored to their teaching style and classroom context. It can aid in adapting instructional materials to meet specific classroom needs, such as ensuring compliance with Individualized Education Programs (IEPs) for students. For instance, AI could facilitate the adaptation of standardized resources to accommodate diverse learning requirements—for example, providing a voice assistant for a visually impaired student or enabling students to present using American Sign Language with AI-assisted translation into spoken English for other students.

Amidst the dynamic classroom environment, teachers may sometimes struggle to fully grasp students' expressions or discern the underlying factors influencing their behaviors. Contextual awareness is key, as teachers may not always be privy to external factors impacting student interactions.

5.4 Designing, Selecting, and Evaluating AI Tools. Currently, educators actively contribute to designing and choosing technologies. They assess usability and practicality, evaluate evidence of effectiveness, and disseminate their findings to school administrators.

5.5 Challenge: Balancing Human and Computer Decision-Making. A significant new challenge with AI-enabled tools for teachers is that these tools can perform tasks autonomously. When teachers delegate work to an AI system, the system may continue independently. Important questions have been raised about managing control in a hybrid teaching environment: When should a teacher maintain control? What tasks can be delegated to a computational system? How can teachers effectively monitor the AI system, override its decisions, or reclaim control when needed?

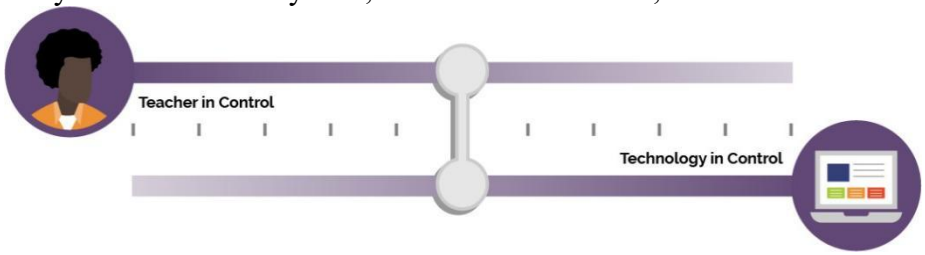

Figure 8. The tension between human and AI decision making: Who is in control?

Figure 8 illustrates the tension regarding control in the classroom. On the left, the teacher has complete control, with no AI involvement. On the right, technology has full control, excluding the teacher—a rarely desirable scenario.

The middle ground is complex and involves numerous choices. Having analyzed various products and proposed some possibilities for finding this balance:

1. The technology only offers information and recommendations to the teacher
2. 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.
3. The teacher delegates more broadly to the technology, with clear protocols for alerts, for monitoring, and for when the teacher takes back control.

It is crucial to identify instructional decisions with varying levels of consequences for students and exercise caution when delegating control over highly consequential decisions, such as course placement or disciplinary referrals. To fully integrate human oversight, AI technologies must enable teacher monitoring, provide protocols to alert teachers when their input is needed, and allow for overrides at the classroom, school, or district level if there is disagreement with an instructional choice.

5.6 Challenge: Making Teaching Jobs Easier While Avoiding Surveillance. We also acknowledge that the technologies designed to simplify tasks can also introduce new possibilities for surveillance (Figure 9). For example, a voice assistant in the kitchen might help with tasks like setting a cooking timer, but it could also pick up on private conversations. This kind of dilemma will also arise in classrooms. When teachers use an AI assistant to record data on their speech, search for teaching resources, or monitor other behaviors, the data can personalize resources and recommendations for them. However, the same data could be used to monitor the teachers, potentially leading to adverse consequences.

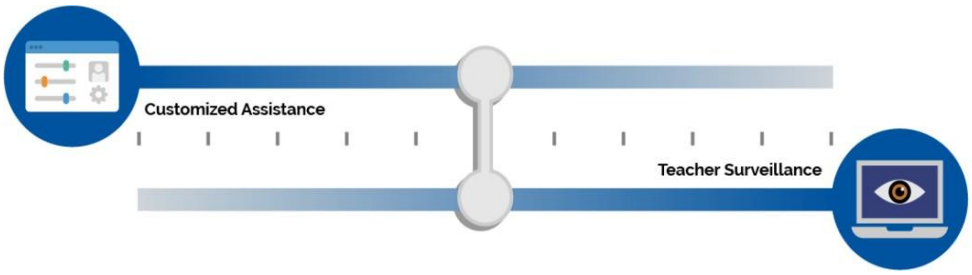

Figure 9. Highly customized assistance vs. increased teacher surveillance

5.7 Challenge: Responding to Students' Strengths While Protecting Their Privacy. In culturally responsive and culturally sustaining approaches, educators design materials to leverage the "assets"—individual, community, and cultural strengths that students bring to their learning. While focusing on these assets, educators must also meet students where they are, recognizing both their strengths and needs. AI could assist by helping teachers customize curricular resources. However, for AI to be effective, the data input into the system must provide comprehensive information about the students. This information could include demographic details, preferences, outside interests, relationships, or experiences. Educators are deeply concerned about the handling, deletion, and visibility of this data. When considering the use of AI-enabled technologies to address educational inequities, educators must ensure that student information shared with or stored in an AI system complies with federal or state privacy laws, such as FERPA.

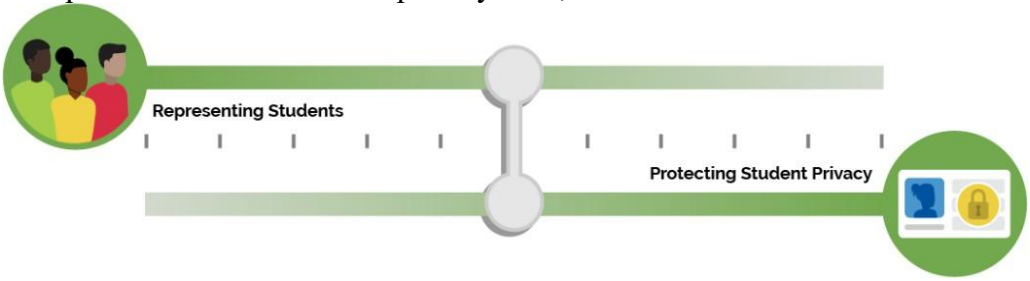

Figure 10. Responding to students' strengths while fully protecting student privacy

AI can introduce algorithmic discrimination through biases present in the data, code, or models within AI-enhanced edtech. Engineers develop AI models using existing data, which may not be representative or might contain associations that contradict policy goals. Moreover, the automations created by engineers for AI might not address the needs of all student groups within a diverse population.

AI has the potential to help teachers customize and personalize materials for their students by leveraging the teacher's understanding of their needs and strengths. Customizing curricular resources is time-consuming, and teachers are already exploring how AI chatbots can assist in designing additional resources for their students.

5.8 Key Recommendation: Inspectable, Explainable, Overridable AI. In the Introduction, we discuss that when AI is integrated into a system, its core component is a model. Some AI models can recognize patterns and take appropriate actions but cannot explain their reasoning, such as how they connected the pattern to the action. This lack of explainability is insufficient for teaching; teachers need to understand how an AI model analyzed a student's work and why it recommended a specific tutorial, resource, or next step.

Explainability of an AI system's decisions is crucial for teachers to evaluate automated decisions. It helps teachers develop appropriate levels of trust and distrust in AI, particularly in identifying where the AI model might make poor decisions.

Teachers need to be able to inspect what an AI model is doing. For example, they need to see which instructional recommendations are being made and to which students, identify students stuck in a loop of remedial work, and recognize those making progress.

6. Formative Assessment

Formative assessment is traditionally a key use of edtech because feedback loops are vital to improving teaching and learning.

6.1 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. 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.

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 have the potential to enhance formative assessments. For instance, AI algorithms can analyze question types that require students to draw a graph or create a model, grouping similar student responses for the teacher to review. AI can also provide learners

with feedback on complex skills, such as learning American Sign Language, speaking a foreign language, or other practice situations where immediate human feedback is unavailable.

In general, an AI assistant can help reduce the workload for teachers by handling the grading of simpler aspects of student responses.

6.2 Implications for Teaching and Learning. Real-time instructional feedback can be valuable for enhancing learning when it assists both learners and teachers in making improvements. However, many students and teachers often have negative experiences with assessments, creating a challenging balance between the advantages of data from formative assessments and the practicalities of implementing additional assessments in educational settings.

Some AI-enabled systems aim to address this issue. For instance, an AI-powered reading tutor listens to students as they read aloud and provides immediate feedback to help improve their reading skills. Students reportedly enjoyed this interactive approach, which also proved effective. Similarly, researchers have integrated formative assessments into educational games, allowing students to demonstrate their understanding of Newtonian physics as they advance through increasingly difficult game levels. Ensuring that students feel safe, confident, and trust the feedback from these AI-enabled systems is crucial for highlighting their learning progress.

AI-enhanced formative assessments also have the potential to save teachers time on tasks like grading, enabling them to spend more time directly supporting and engaging with students.

6.3 Insight: AI Can Enhance Feedback Loops. The term "formative assessment" goes beyond merely referring to a test or measurement. An assessment becomes formative when it leads to meaningful reflections and adjustments in teaching, learning, or both. The concept of "feedback loops" highlights that measurement is just a part of the process. Effective feedback loops, which result in instructional improvements and adaptations in both teaching and learning, produce the most significant outcomes for students.

Students benefit from feedback when they are working individually, participating in small groups, or engaging in classroom discussions. Feedback loops are particularly valuable "in the moment" as a student practices a skill, but they are also crucial over longer periods, such as after completing a project or term paper. Additionally, feedback loops can aid teachers by helping them recognize their own patterns in responding to students' ideas.

6.4 An Example: Automated Essay Scoring. An illustrative case is Automated Essay Scoring (AES). Developing strong writing skills is crucial for students as it prepares them for success in college and career pursuits. However, providing detailed feedback on essays is time-consuming for human reviewers. Ellis Page envisioned computer programs capable of evaluating and offering feedback on student essays as early as 1966, and significant advancements have been made in AES technologies over the past 56 years. Today, readers are likely familiar with applications such as Grammarly, Turnitin, and various essay analysis engines used by publishers and assessment companies.

Writing proficiency is fundamental to academic and professional aspirations, and aspiring writers benefit greatly from comprehensive feedback.

AI-driven AES technologies can analyze certain aspects of student essays but may be misled by factors like essay length or instances where students strategically place keywords in sentences that lack coherence—a flaw easily identified by human readers.

6.5 Key Recommendation: Harness Assessment Expertise to Reduce Bias. Bias and fairness are critical considerations in assessment design and implementation, particularly within the realm of AI-enabled assessment. In traditional assessments, bias may occur if test items contain unnecessary details that advantage certain groups over others—for example, a story-based question referencing a sport predominantly played by boys may disadvantage girls.

Algorithmic discrimination extends beyond the measurement aspect of formative assessment; it also encompasses the feedback loop and the instructional interventions triggered by data collected through assessments. When algorithms recommend hints, next steps, or resources to students, it's crucial to ensure fairness—ensuring all groups receive equally useful assistance and avoiding discriminatory outcomes.

Fairness considerations in AI-enabled formative assessment involve not only providing opportunities for learning through feedback loops but also ensuring the quality of learning within and beyond these loops.

7. Research and Development

Policy development hinges on evidence-based research, just as refining practice relies on iterative feedback loops driven by empirical evidence.

7.1 Insight: Research Can Strengthen the Role of Context in AI. Despite the significance of the grand challenges outlined in the 2010s, there is a noticeable shift in the R&D community's focus today. While the challenges from the previous decade primarily addressed technical issues, contemporary researchers are increasingly emphasizing the exploration of context. Similarly, tech companies are striving to develop platforms that can adapt to learners' characteristics and broader situations, moving beyond narrow cognitive attributes.

Researchers are expanding the concept of "adaptivity" to include support for group learning, a prevalent form of learning in schools across the U.S. The emphasis on context is not incidental; it has long been recognized as a fundamental challenge in AI.

However, agreeing to prioritize context poses challenges. As depicted in Figure 11, there exists a tension between the depth of context and the rapid pace of technological advancements in AI R&D. While AI is often driven by the urgency to innovate and scale new applications quickly, researchers and developers acknowledge that achieving meaningful innovations in AI education necessitates integrating context into the process early and consistently.

In conclusion, embedding context early and persistently in AI R&D is crucial. It ensures that technological advancements in education align closely with educational objectives and effectively address the complexities of learning environments.

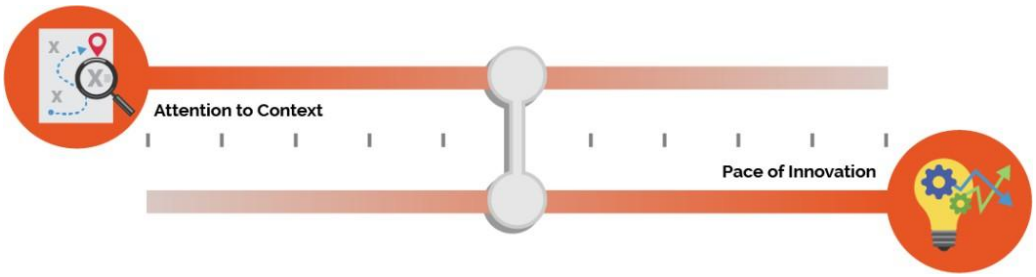

Figure 11. The tension between depth of context and pace of technological advances in AI

Moreover, an increased emphasis on context within this field will significantly alter the landscape of R&D. Contextual considerations are multifaceted, and this shift will prompt various types of changes in R&D practices.

Below, we outline these four types and subsequently provide detailed exploration of each, highlighting their significance as catalysts for rethinking R&D approaches.

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?

7.2 Connecting with Public Policy. Defining human-centered AI for education necessitates adopting a foundational principle centered on human values. This approach should guide the development of policies governing the broader application of AI across society. Issues such as power dynamics between companies and consumers, particularly regarding data ownership, are mirrored in the educational ecosystem. Simultaneously, public discourse on ethics, bias, responsibility, and related concepts occurs both in public policy and educational contexts.

A key insight is the potential importance of enhancing AI literacy in education for broader societal benefit. For instance, panelists highlighted the necessity of improving AI literacy to influence the design of these technologies. Another researcher emphasized creating educational technology environments that foster open dialogue across diverse perspectives, reflecting current societal challenges.

Researchers are advocating for a shift beyond superficial approaches to tackling these issues seriously. They recognize the need to engage more deeply with policy, an area where many edtech and AI researchers may lack extensive experience. Developers, while familiar with certain policy issues such as data privacy and security, are increasingly involved in discussions about ethics, bias, transparency, and other critical topics.

8. 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?

8.1 Insight: Aligning AI to Policy Objectives. Each section of this policy report offers insights into four key areas: learning, teaching, assessment, and research.

- To support our goal of attracting and retaining teachers, a nation should prioritize AI assistants designed to enhance teaching roles and equip educators with essential information for fostering close, empathetic connections with students. Emphasizing the involvement of teachers ensures that AI-enabled classroom technologies keep educators informed, engaged with their students, and empowered to make critical instructional decisions. Maintaining teacher involvement in these processes is crucial for managing risks effectively.
- To promote equitable learning, particularly for those disproportionately affected by the pandemic, AI has the potential to transform educational technology from its current deficit-oriented approach to one that focuses on strengths. Instead of solely identifying student weaknesses and providing remedial measures, AI-powered edtech could recommend strategies that leverage students' existing strengths across cognitive, social, and self-regulatory domains, thereby enhancing learning outcomes significantly. This holistic approach should encompass support for students with disabilities and English learners to ensure inclusivity and equity. Addressing equity requires a vigilant awareness of inherent biases in AI system development and decisive measures to uphold fairness.
- Regarding pathways to successful careers, AI-enabled assessments have the potential to offer students and teachers formative guidance on a diverse array of essential skills, aiming to enrich the learning experience with valuable information. Embracing a human-centric perspective, we should adopt a systemic approach to assessments, ensuring that students, teachers, and other stakeholders remain central to instructional decision-making processes.
- In the realm of equity, as research progresses and incorporates more contextual understanding into AI, we can leverage AI to support objectives that demand personalized learning resources. This includes empowering teachers to efficiently adapt materials to accommodate neurodiverse learners and enhance responsiveness to the unique needs of local communities and diverse cultures.

Moving forward, educational leaders must consistently bring these policy priorities alongside their own to every AI discussion, steering the dialogue toward human-centered concerns rather than solely focusing on the potential capabilities of new technologies.

8.2 Calling Education Leaders to Action. We summarize recommendations for policy action. These recommendations are for education leaders.

8.3 Recommendation #1: Align AI Models to a Shared Vision for Education. Throughout this report, we have underscored that AI technologies are built on models that inevitably have limitations. It falls upon humans to define educational objectives and assess how well these models align with and contribute to those goals—or conversely, how they may fall short and potentially pose harm. While this evaluation of tools against educational priorities may seem straightforward, the allure of technology can sometimes lead to a mindset of "let's explore what the tech can do," diverting attention from critical educational objectives and potentially leading to the adoption of ill-fitting models.

Therefore, we urge educational policymakers and decision-makers at all levels—local, state, and federal—to wield their influence in aligning priorities, educational strategies, and technology adoption decisions with the paramount educational needs of students. Local education leaders are best positioned to understand their urgent educational priorities.

Every discourse on AI or any emerging technology must commence with a deliberate emphasis on the educational needs and priorities of students, concluding with an evaluation of effectiveness centered on those same needs and priorities. Equity, notably, remains a critical priority that demands continual vigilance, particularly given concerns over biased AI models.

We specifically urge leaders to refrain from over-idealizing AI or fixating solely on promising applications and outcomes. Instead, we encourage a rigorous inquiry into how AI-enabled systems and tools function within educational settings.

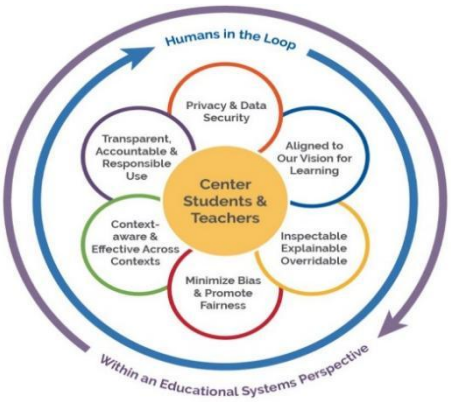

Figure 12. Recommendation for desired qualities of AI tools and systems in education

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.

Teachers, students, and their families contribute significantly to adoption decisions. It is essential for leaders and parents to support educators when they question or override an AI model based on their professional expertise. Additionally, technology developers must be transparent about the models they employ. Policymakers should consider mandates for disclosure to ensure the marketplace operates on informed decisions about AI models rather than solely on marketing claims.

Furthermore, government oversight is essential. AI models are human-created approximations of reality and should be subject to policies requiring transparency, inspectability, explainability, and override capability within educational systems.

8.4 Recommendation #2: Develop Education-Specific Guidelines and Guardrails. The final recommendation focuses on policymakers. The country's educational system places a significant emphasis on local decision-making, but as technology rapidly advances in complexity, local leaders face challenges in making well-informed decisions regarding the deployment of artificial intelligence. Beyond traditional concerns such as data privacy and security, new issues such as bias, transparency, and accountability also come into play.

Regulations concerning crucial 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) need reassessment and further exploration in light of emerging technologies in educational settings. Similarly, laws like the Individuals with Disabilities Education Act (IDEA) may require updates as AI-enabled learning technologies evolve.

The Department encourages coordinated efforts across all levels of the education system. In addition to federal laws, many states have enacted privacy legislation governing the use of educational technology and platforms in schools. While general frameworks for responsible AI in sectors like health and consumer products can offer insights, they may not fully address the specific needs of education. Leaders at every level must be prepared to address issues beyond privacy and security, including potential biases and fairness, as they navigate the complexities of AI in education.

Acknowledgements

Project Team

Artificial Intelligence in Education was developed under the leadership and guidance of Wang Zhong Sheng, Supervisor Xu Fei of the Office of Educational Technology.

Creation of this document was provided by Tatenda Mavhunga, Special thanks to Research team in Algorithm Design team in Computer Science in Xi'an Technological University.

References

[1] S. Akgun, C. Greenhow: *AI Ethics*, Vol. 2 (2022), p.431–440.

[2] V. Aleven, E. A. McLaughlin, R. A. Glenn and K. R. Koedinger: *Handbook of research on learning and instruction*, (2016), p.522-560.

[3] R.S. Baker, L. Esbenshade and J. Vitale, S. Karumbaiah, (2022).

[4] P. Black, D. WiliamP: *Phi Delta Kappan,* Vol.92 (1998), p.81-90.

[5] P. Black, D. Wiliam: *Educational Assessment, Evaluation and Accountability*, Vol. 21 (2009) p. 5-31.

[6] M.A. Boden: *Artificial intelligence: A very short introduction.*

[7] I. Celik, M. Dindar, H. Muukkonen and S. Järvelä: *TechTrends*, Vol. 66 (2022), p.616–630.

[8] C. Chen, H.W. Park and C. Breazeal: *Computers & Education*, (2020).

[9] G. Chen, S. Clarke and L.B. Resnick: *Technology, Instruction, Cognition and Learning*, Vol.10 (2015), p.85-105.

[10] E. Dieterle, C. Dede and M. Walker: *AI & Society*, (2022).

[11] A. Doewes, M. Pechenizkiy: *Proceedings of the 14th International Conference on Educational Data Mining*, (2021).

[12] D.C. Englebart: *Augmenting human intellect: A conceptual framework.*

[13] Z. Ersozlu, S. Ledger, A. Ersozlu, F. Mayne and H. Wildy: *SAGE Open*, (2021).

[14] S. Forsyth, B. Dalton, E.H. Foster, B. Walsh, J. Smilack and T. Yeh: *2021 Conference on Research in Equitable and Sustained Participation in Engineering, Computing, and Technology* (RESPECT), (2021).

[15] L. Friedman, N. Blair Black, E. Walker, and J. Roschelle: *Association of Computing Machinery BLOG@*, (2021).