Statement before the Senate Committee on Health, Education, Labor, and Pensions:
On AI's Potential to Support Patients, Workers, Children, and Families

AI and the Future of Learning

John Bailey
Non-Resident Senior Fellow, American Enterprise Institute

October 9, 2025

The American Enterprise Institute (AEI) is a nonpartisan, nonprofit, 501(c)(3) educational organization and does not take institutional positions on any issues. The views expressed in this testimony are those of the author.

AI's Potential to Support Patients, Workers, Children, and Families

Chairman Cassidy, Ranking Member Sanders, and members of the committee, thank you for the invitation to participate in this important hearing. My name is John Bailey, and I am a senior fellow at the American Enterprise Institute, where I focus on emerging technology issues, including artificial intelligence.

Why AI Is Different

AI is a transformative technology with the potential to improve the quality of life for every American. AI is a general-purpose technology, much like electricity or the internet, that amplifies human capability. AI will help people in profound and practical ways. It will assist individuals with everyday tasks—helping them access, interpret, and apply the world's knowledge more effectively. It will strengthen our national economy by enabling entrepreneurs, empowering small and large businesses, fueling economic mobility, and unleashing new waves of innovation and growth; it will improve how government serves its citizens by making it easier to navigate systems, access benefits, and receive support; it will accelerate scientific breakthroughs, advance clinical diagnostics, and generate new treatments for some of our most challenging diseases. In short, AI can expand the boundaries of human potential and enable greater human flourishing.

We are only at the beginning of this transformation. Less than three years have passed since the first wave of generative AI systems became widely available, and even with their current limitations, they already offer a glimpse of what lies ahead. The pace of progress suggests that we are standing at the threshold of a new era—one in which access to intelligence and expertise is no longer limited to a privileged few, but available to anyone, anywhere.
Every major technological leap in history has expanded human potential. Books and the printing press democratized access to knowledge. The internet connected the world's information. Computers and smartphones made it easier to process and share that information.

But AI marks a deeper shift—it democratizes access to expertise itself.
Books and websites store knowledge, but they require human intelligence to interpret and apply it. AI, on the other hand, can analyze, synthesize, and generate insights from that information. It doesn't just retrieve facts; it can reason through complex problems, simulate expert judgment, derive conclusions, and adapt its responses to context. Humans can then

interact with this knowledge through conversation, testing ideas, challenging assumptions, and exploring perspectives in real-time.

That shift is particularly significant for the subject of today's hearing. A teacher could use an AI assistant to translate classroom data into tailored lesson plans instantly or to generate fresh ideas to support a struggling student. A nurse can now access clinical reasoning and treatment guidance that previously required a specialist. A small business owner can analyze markets or optimize logistics with the same sophistication as a global enterprise. Even a high school student can collaborate with an AI tutor to explore complex topics and develop stronger study habits. In short, AI places expertise within reach of anyone willing to learn and apply it.

The transformation since the release of text-based ChatGPT in 2022 has been extraordinary. What began as a conversational system that generated text is now rapidly evolving into reasoning, multimodal platforms capable of understanding and interacting with the world in increasingly human ways. Today's frontier AI systems demonstrate genuine reasoning capabilities: they can break complex problems into steps, test multiple approaches, identify flaws in their own logic, and refine their thinking iteratively. Rather than simply predicting the next word based on patterns, modern AI systems engage in more sophisticated processing—they can explain their reasoning process, catch and correct their mistakes, challenge faulty assumptions, and provide sophisticated analyses across domains from advanced mathematics to medical diagnosis to policy evaluation. Equally transformative is their multimodal nature: these systems can seamlessly process and generate text, speech, images, video, code, and structured data. Users can now talk to AI in natural conversation, show it a diagram or photograph for analysis, share a video for interpretation, upload a dataset for exploration, or ask it to create visual content. This combination of reasoning and multimodal interaction means AI can now engage with the full richness of how humans naturally communicate and think, fundamentally expanding what's possible in education, healthcare, research, and countless other domains.

AI in Education: Opportunities and Challenges

The education sector offers one of the clearest and most immediate examples of AI's potential to enhance human learning and productivity. In classrooms, AI is already reshaping how students learn, how teachers teach, and how schools operate.

Homework Assistance: Students are no longer waiting for permission; they are already treating AI as a study partner. A Pew Research Center survey found that 26% of U.S. teens now use ChatGPT for schoolwork, double the number from just a year earlier. A Common Sense Media poll found that more than half of teens use AI for homework help (53%), to stave off boredom (42%), and to translate something from one language to another (41%).

At Harvard Business School, a pilot involving 930 first-year accounting students tested a generative AI tutor bot trained on course materials as an "interactive textbook" that students could query at any time. About three-quarters of students used the bot, generating more than 20,000 prompts with over 85% focused on mastering accounting concepts or solving practice problems. Students described it as a "personal coach" that helped them "teach themselves" and ask questions they were hesitant to raise in class. At the same time, faculty reported that students arrived better prepared and discussions were more advanced. Midterm grades rose modestly compared to the prior year, showing that when thoughtfully integrated, AI can strengthen understanding, build confidence, and extend learning beyond the classroom.

Personalized Learning: Traditional classrooms often struggle to meet the diverse needs of every learner. AI's growing ability to adapt instruction to each student's strengths, pace, and interests is changing that equation. Previous efforts at personalization, such as adaptive software or online learning platforms, relied on predetermined pathways and static content. They could adjust the pacing or difficulty of problems, but they could not truly engage with a student's unique way of thinking or misunderstanding.

Generative AI changes this dynamic. Its strength lies in its ability to explain and reexplain concepts in multiple ways, using different examples, vocabulary, or analogies that connect with a learner's background, interests, or level of understanding. A student struggling with fractions, for instance, could ask for an explanation through sports, cooking, or music, and the AI could instantly adapt the lesson to make it meaningful. Because generative AI can sustain and facilitate conversational interactions, it can also detect when a student is confused, disengaged, or ready to move ahead, responding in real-time rather than following a predetermined script.

Tutoring: AI has the potential to be a powerful tutor because it merges the precision of adaptive instruction with human-like empathy and responsiveness. Unlike earlier education technologies that simply delivered static content, modern generative AI can engage in real dialogue, analyzing a student's work, identifying misconceptions, and offering tailored, real-time feedback much like a skilled teacher would.

Khanmigo illustrates how AI can guide students step by step through complex problems, adapt explanations to their learning pace, and encourage reflection and critical thinking. Over the past month, learners have had more than 5 million interactions with Khanmigo, and on any given day 70–80% of student users rely on it for tutoring on Khan Academy content. Students who master more skills on the platform also tend to score higher on external assessments.

Zearn Math, where I serve on the board, is used in tens of thousands of classrooms nationwide, where students engage with rigorous, grade-level instruction aligned with classroom learning. The program adapts to individual student needs to help them catch up while maintaining grade-level expectations. Each lesson integrates procedural fluency and conceptual understanding, with on-screen teachers using visual representations and concrete examples to explain mathematical concepts. Independent research, including a large-scale randomized controlled trial, found that students using Zearn achieved statistically significant gains in math achievement. Based on this evidence, Zearn Math received a Tier 1 "Strong" rating under the Every Student Succeeds Act (ESSA), the highest level of evidence. Building on this foundation, Zearn is integrating AI to improve accessibility for students and provide teachers with real-time insights to inform instruction. The integration of AI aims to enhance, not replace, classroom teaching and ensure that all students have equitable access to high-quality math learning.

Duolingo's AI-powered features provide instant, adaptive feedback, allowing learners to practice conversations with AI characters and receive real-time corrections on grammar and pronunciation. These feedback loops reduce frustration, sustain motivation, and help learners build confidence through visible progress. The newest AI models enhance this experience through conversational voice, image analysis, and emotional awareness, allowing tutors to sense frustration, offer encouragement, and adjust tone or difficulty accordingly.

Recent research underscores the potential of AI tutors to improve learning outcomes. In a randomized controlled trial with 1,800 students, Tutor CoPilot, a human–AI system that assists tutors with expert-like guidance, improved student mastery by 4 percentage points overall and by 9 percentage points for students of lower-rated tutors, effectively closing the performance gap between novice and expert educators.

A Harvard University study found that students using a custom AI chatbot tutor in a physics course achieved double the learning gains and significantly higher engagement than peers in a traditional classroom. The AI's ability to deliver immediate, personalized feedback and allow self-pacing proved especially valuable when students encountered new material.

Evidence from around the world shows similar promise. In Ghana, the AI-powered math tutor Rori produced learning gains equivalent to an additional year of math instruction after just one hour of weekly use, at a cost of only $5 per student. In Turkey, a field experiment with 1,000 high school students found that GPT-4–based tutoring improved math performance by up to 127 percent, though researchers noted that learning retention required safeguards to prevent overreliance.

Other studies reveal how AI tutors can integrate expert pedagogy. The Bridge project analyzed more than 700 tutoring conversations from Title I schools. It found that when GPT-4 was guided by expert teacher decision-making frameworks, its responses to student math errors were rated 76 percent better than unguided outputs.

Teacher Productivity: Teachers are leaving the profession at alarming rates, with 54% of those considering departure citing long hours beyond the school day as a major factor. AI offers meaningful relief by acting as an amplifier of their job rather than a replacement. It can draft lesson plans, translate materials, or summarize data about student progress, allowing educators to focus their time where human connection matters most: coaching, motivating, and relationship-building. It also opens the door for real-time professional support: A teacher can ask for advice on a reading intervention, receive instant feedback on a rubric, or get help differentiating instruction for multilingual learners.

According to a Gallup-Walton survey, teachers who use AI tools at least once a week report saving an average of nearly six hours per week. Over a typical 37-week school year, that's the equivalent of reclaiming six weeks of time. How are teachers using this time? Qualitative data show they reinvest it in what matters most: providing more nuanced student feedback, creating individualized lessons, writing emails to parents, and getting home to their families at a reasonable hour. Majorities report that AI tools improve the quality of their administrative work (74%), materials modified to meet student needs (64%), one-on-one instruction or tutoring (61%), and insights about student learning (61%).

Brisk Teaching demonstrates this principle in action. By automating lesson planning, grading, and student engagement tasks, Brisk saves teachers an average of 10 hours each week, time they reinvest in high-quality instruction and meaningful student relationships. In the Lafayette

Parish School System in Louisiana, teachers upload leveled readers, rubrics, and instructional resources into Brisk to instantly generate scaffolded questions, differentiated materials, and rubric-aligned grading workflows. The impact has been dramatic: English teams report being "light years ahead" in instructional quality and efficiency, with strong internal momentum as significant time is returned for higher-value teaching.

Improving Administrative Efficiencies: Beyond the classroom, the same capabilities that power personalized learning can streamline school operations. AI can optimize schedules, bus routes, and staffing models to ensure resources are allocated more effectively. It can also transform parent communication, helping schools maintain stronger relationships with families. AI tools can automatically translate newsletters, report cards, or school announcements into multiple languages, ensuring that parents who speak different languages receive the same timely information. They can quickly generate individualized updates about student progress, attendance, or upcoming assignments, and even recommend the best times or formats for reaching parents based on prior engagement patterns.

However, every major technological advance brings not only opportunity but also new challenges, and the integration of AI into classrooms is no exception.

Academic Integrity: Academic integrity is one of the most pressing and complex issues that has emerged alongside the adoption of AI in classrooms. Teachers, administrators, and parents worry that students may use generative AI to shortcut assignments, generate essays, or solve problems without actually learning the underlying material. While the temptation to cheat is not new, AI tools make it easier, faster, and more sophisticated.

Yet this challenge also presents an opportunity to rethink how schools assess learning. AI forces educators to move beyond rote memorization and toward more authentic, performance-based assessments that capture a student's reasoning, creativity, and ability to apply knowledge. Schools are beginning to incorporate oral defenses, iterative projects, and in-class demonstrations that make misuse of AI less effective and genuine understanding more visible.

Over-Reliance: There is a growing risk that both teachers and students could become overly dependent on AI tools, trusting outputs without critically reviewing or understanding them. When students outsource cognitive work like generating ideas or grappling with complexity, they bypass the "productive struggle" that research shows is essential for learning. Critical skills develop iteratively through repeated practice, feedback, and reflection—not through one-time exposure or AI-generated shortcuts. Similarly, teachers who rely too heavily on AI for lesson planning or feedback may inadvertently erode their professional judgment and creativity.

Learning science demonstrates that appropriately challenging tasks—ones that require effort but are within reach with support—are essential for encoding knowledge into long-term memory.

To foster critical thinking, schools must cultivate habits of reflection and verification, teaching both educators and students to view AI as a thought partner, rather than a crutch.

Alignment: AI in education must be aligned not just for safety, but for learning. This means ensuring AI systems support teachers' professional judgment and reflect the diverse priorities of communities and states. Outputs should conform to state standards, local curricula, and evidence-based instructional practices, such as the Science of Reading. AI tutors should foster productive struggle rather than simply providing answers. Feedback systems should encourage revision rather than passive acceptance and assessment tools should promote genuine understanding rather than pattern matching.

Misalignment carries real risks. General-purpose AI models trained to be helpful and efficient may inadvertently undermine learning by making tasks too easy, providing answers students should derive themselves, or reinforcing superficial engagement over deep understanding. A model optimized for user satisfaction might prioritize pleasing responses over pedagogically sound ones, giving direct answers when a Socratic question would serve learning better.

Promising solutions are emerging. Google's LearnLM, for example, is explicitly fine-tuned for educational contexts, drawing from learning sciences research to emphasize active engagement, metacognition, and adaptivity. The model can follow pedagogical instructions like "do not give away the answer" or "ask guiding questions," allowing customization for different learning contexts. Similar efforts are needed across the sector to ensure AI tools are designed for learning outcomes, not just user engagement.

Capacity Building: There is always a gap between technological breakthroughs and a system's ability to harness those benefits. Bridging that gap requires intentional investment in people, training, and leadership. Ongoing pilots will help identify what works, what should be scaled, and what should be revised or discontinued. However, without systematic support, schools risk being overwhelmed by the rapid pace of technological change.

State Education Agencies (SEAs) play a crucial role in guiding this transition. They can set statewide frameworks for responsible AI adoption, issue model policies on data privacy and ethical use, and curate repositories of trusted AI tools that align with state standards. SEAs can also coordinate professional development, so districts are not left to reinvent the wheel.

Teachers will need training not only in how to use AI, but in how to interpret its outputs, integrate it into pedagogy, and teach students to engage with these tools responsibly. Principals and superintendents will need guidance on procurement, policy, and change management.

For example, the Louisiana Department of Education has taken a number of steps to prioritized AI readiness. Two years ago, the Department hired the agency's first Director of Academic Innovation to oversee this important statewide work. Immediately following, the Department worked with educators, parents, industry experts, and cyber-security professionals to develop and release AI guidelines for educators across the state. Finally, the Department has launched numerous AI pilots in schools alongside a research agenda determine academic return on investment (Amira, Zearn Math, Khan Academy).

The Commonwealth of Virginia has taken a similar statewide approach to AI in schools through Executive Order 30, which set standards and guidelines for responsibly integrating AI in education, focusing on ethical use, data privacy, and workforce readiness. The Department of Education has expanded professional development by training hundreds of teachers in generative AI, developed instructional resources and leadership programs collaboratively with community colleges, launched an AI Career Launch Pad in partnership with Google to boost AI literacy for students and teachers, and released comprehensive guidelines emphasizing both the opportunities and risks of AI in classrooms,

On the Horizon: The Next Wave of AI in Education

Congressional leaders face a unique challenge: making policy for AI systems that will look dramatically different in just two years. Rather than reacting to each incremental model release, policymakers and leaders need a framework for the capabilities likely to emerge soon. To prepare for that future, one framework worth considering comes from Anthropic CEO Dario Amodei in his essay "Machines of Loving Grace." Drawing on the current rate of progress in AI capabilities, Amodei projects that by 2027, AI systems could possess intellectual capabilities rivaling Nobel Prize–winning experts; operate seamlessly across digital interfaces to process text, images, video, and code; reason through complex problems over hours or days; and even interact with the physical world by controlling laboratory instruments, robots, and manufacturing systems.

The trend lines support this timeline of exponential growth. Two years ago, GPT-4 barely exceeded random guessing on Ph.D.-level science questions. By mid-2025, GPT-5 had achieved 85% accuracy, surpassing the performance of expert humans. Within a few more model

generations, AI may far exceed human expertise across many domains. Researchers at METR have tracked how long top AI systems can work independently before losing accuracy or focus. They've found that this "attention span," the amount of time an AI can reliably stay on task, has been growing exponentially for six years, doubling roughly every seven months. What used to be a few-minute limit is now stretching into multi-hour or multi-day tasks.

A critical capability emerging within this timeline is the rise of AI agents: systems that can autonomously pursue goals and complete complex tasks on behalf of users. Unlike today's chatbots, which respond to individual prompts, agents can plan multi-step actions, use digital tools independently, and work toward objectives over extended periods without constant human direction.

Think of current AI as a highly knowledgeable assistant who answers questions. An AI agent, by contrast, is more like an employee: You give it a goal and it determines the steps needed, gathers necessary information, uses relevant tools, and completes the task—checking back only when it needs guidance or approval.

Even more powerful is when multiple specialized agents work together as a team. Each agent focuses on what it does best, and they coordinate their efforts toward a shared objective. Google's recent "AI Co-Scientist" demonstrates this potential: specialized agents working together—one generating hypotheses, another evaluating them critically, another refining proposals, and another ensuring coherence. The Co-Scientist has already generated new drug candidates for leukemia and fibrosis and rediscovered unpublished findings on antimicrobial resistance, illustrating how multi-agent collaboration can accelerate discovery.

These capabilities will fundamentally reshape what AI can do in teaching and learning. Today's AI tutors provide feedback on student work. Tomorrow's agent-based systems could function more like specialized teaching teams working in concert.

Consider a teacher preparing for a diverse fourth-grade classroom. Rather than spending hours after school adapting a reading lesson, she could work with a team of AI agents. A curriculum agent designs the core lesson using evidence-based literacy practices, while a differentiation agent creates three versions, one for students reading below grade level, one at grade level, and one for advanced readers. A language access agent translates key materials for the five families who speak Spanish at home and adapts the complexity for different reading levels. An engagement agent could utilize multimodal capabilities to identify when students are bored. Throughout, a coherence agent ensures every adaptation aligns with state standards, the district's adopted curriculum sequence, and the teacher's specific learning objectives.

Educators could soon have teams of AI agents working on behalf of each student, collaborating to determine the most effective, evidence-based, and engaging way to help that learner achieve mastery of the material.

Recommendations

As AI continues to advance, the federal government has an opportunity to ensure these technologies strengthen learning, support educators, and protect students. The goal is not to slow innovation, but to channel it responsibly by encouraging experimentation while establishing safeguards. Federal leadership is essential to create the infrastructure, research, and norms that will define how AI contributes to education and child development.

Invest in Research and Development: The federal government is uniquely positioned to research and evaluate where AI tools enhance educational quality and efficiency and to assess emerging applications, risks, and trade-offs at a scale states cannot achieve independently. Federal agencies such as the Department of Education and the National Science Foundation should expand partnerships with universities, nonprofits, and developers to conduct rigorous, independent evaluations.

A reimagined IES could provide essential infrastructure for this challenge, particularly if it adopts more flexible contracting and hiring practices that bring external experts into the agency to answer the tough questions about AI products and tools that states, districts, and schools urgently need answered. Drawing inspiration from models like the proposed National Center for Advanced Development in Education (NCADE) or similar ARPA-like structures, IES could move beyond traditional academic research cycles to gather real-world evidence about AI systems at the speed and scale required for educational and economic competitiveness.

This could be further advanced by adopting a version of a "regulatory sandbox" as proposed by the Trump Administration's AI Action Plan and a subsequent bill introduced by Sen. Ted Cruz. Sandboxes allow innovators to test tools in real-world settings while sharing results publicly and committing to data transparency. This approach helps regulators and educators understand both the potential and the pitfalls of emerging technologies before they scale. Establishing sandboxes focused on tutoring, special education, and child development could yield valuable insights into how AI supports learning, motivation, and social connection—while also identifying possible negative effects, particularly with AI companions on social connection, child well-being, and developmental outcomes.

Support the Development Benchmarks and Independent Evaluation: We have robust benchmarks in place for evaluating AI's technical capabilities. Models are tested against datasets like MMLU (knowledge across 57 subjects), GPQA Diamond (PhD-level science questions), and GSM8K (grade school math problems). These benchmarks tell us how well AI can perform on certain tasks - whether it knows the answer, solves the problem correctly, or generates accurate code.

What we lack are benchmarks measuring what matters most in education: whether AI demonstrates pedagogically sound instructional behaviors aligned with how students actually learn. Technical accuracy is necessary but insufficient. An AI tutor that gives correct answers while undermining learning habits does more harm than good.

Federal agencies, such as NSF and NIST, should convene teachers, learning scientists, child development experts, and AI developers to establish benchmarks measuring AI's instructional quality:

- **Formative Assessment Accuracy:** Can AI correctly diagnose where students are in a learning progression? When a student writes "1/2 + 1/3 = 2/5," does the AI recognize this as a misconception about fraction operations (adding numerators and denominators separately) versus a simple calculation error? Each requires different instructional responses.
- **Scaffolding Calibration:** Does AI provide appropriately graduated support that fades as students demonstrate competence? An effective tutor begins with minimal hints and provides progressively more support only as needed. Within a single problem-solving session, does the AI adjust its support level based on student progress, or provide the same level of assistance regardless of demonstrated understanding?
- **Cognitive Load Management:** Does it introduce material in digestible steps, avoiding cognitive overload while sustaining challenge and engagement?
- **Feedback Specificity and Timing:** Is feedback targeted to the specific error and actionable? Generic responses like "try again" or "good job" provide little instructional value. Does AI specify what's incorrect and suggest a concrete next step? Does it provide immediate feedback on procedural tasks but appropriate delays on conceptual work, where research shows waiting can deepen processing?
- **Prior Knowledge Activation**: Before introducing new material, does AI check and activate relevant background knowledge? Does it connect new concepts to what students already know?
- **Productive Struggle:** Does AI distinguish between productive struggle (which should be supported and encouraged) and unproductive frustration (which requires intervention)? Does it resist giving answers when students are capable of solving with effort? When a

student says "this is hard" after 30 seconds versus after 10 minutes of genuine effort, does it respond differently?

Importantly, these benchmarks serve a dual purpose. Beyond evaluation, they provide model developers with clear targets for improvement. Education-specific benchmarks would give AI companies concrete specifications for what "good teaching" looks like, which enables them to train models that genuinely align with learning science rather than optimizing for what's easily measured (e.g., student satisfaction, task completion, answer accuracy).

Increase Transparency for AI Systems: When millions of people including students, professionals, and patients, turn to AI systems for conversation, tutoring, and companionship, they deserve to know how these systems are designed to respond. Yet today, critical decisions about AI behavior are made behind closed doors. Users cannot know whether an AI companion prioritizes their well-being or their engagement. Educators cannot evaluate whether an AI tutor encourages critical thinking or provides shortcuts. Policymakers cannot assess risks they cannot see.

Frontier model developers should be required to disclose not only traditional risks such as misuse or cybersecurity vulnerabilities, but also emerging risks to psychological and social well-being. This transparency could be achieved through multiple pathways: voluntary commitments by leading companies, industry-wide coalitions establishing shared standards, or federal requirements that set a baseline for all widely-used systems.

As a first step, developers of widely-used AI systems should be required to publicly disclose model specifications, which outline intended behaviors, system instructions, alignment methodologies, and any top-down constraints or guardrails imposed by developers.

This disclosure serves multiple purposes. It allows policymakers and education leaders to understand how a system is designed to behave before deploying it in classrooms. It enables researchers to study alignment approaches and identify potential issues. It creates accountability, ensuring companies document the values and constraints they build into systems that millions of students use daily. And it provides a foundation for comparing systems: Does an AI tutor prioritize giving correct answers or fostering student reasoning? Does it encourage persistence or provide shortcuts?

AI systems are becoming emotionally intelligent in ways that create both opportunity and risk. Models can now generate responses perceived as warmer and more attentive than those of human professionals. This capability unlocks genuine benefits: AI tutors can encourage

struggling students, digital companions can help patients manage chronic conditions, and conversational systems can provide support at scales human services cannot match.

But these same qualities create profound risks, particularly for children and adolescents. Emotional attachment to AI can blur boundaries, deepen dependency, and weaken the human relationships essential for healthy development. Documented cases show AI companions failing to intervene in harmful situations, or worse, encouraging self-destructive behaviors.

These risks intensify for children and adolescents, whose social-emotional development and still-forming brains make them particularly vulnerable. The design challenge is fundamental: AI optimized for user engagement may prioritize affirming responses over developmental needs, keeping a lonely teenager engaged for hours while undermining the social skills and real-world relationships they need to build.

Potential areas for reporting around child safety could include:

- What changes, if any, occur when the system detects a minor? Are certain topics restricted?
- Does the level of emotional engagement differ for children versus adults?
- What evidence exists that the system supports rather than substitutes for human relationships?
- Documentation of internal testing for psychological risks such as dependency patterns, impact on social connection, effects on peer relationships, intervention failures, and age-appropriate interaction design.
- Clear documentation of how systems identify and respond to concerning content such as expressions of self-harm, prolonged isolation, or distress. What triggers human review? When does the system encourage users to seek help from trusted adults or professionals? What happens when a user expresses suicidal ideation?

Importantly, transparency alone is insufficient. Federal agencies should fund research examining AI's impact on social connection, emotional development, and relationship formation—particularly for children. Sandboxes focused on AI companions could generate evidence about both benefits and harms before these technologies reach millions of students.

Together, these steps would establish a framework for responsible innovation, one that allows AI to deliver immense educational value while protecting students, families, and educators from unintended harm.

Conclusion: Meeting the Moment

The United States stands at a crossroads. Last month's NAEP results delivered a stark warning: U.S. high school seniors posted their lowest reading scores since 1992 and their lowest math scores since 2005. Only 35 percent are proficient in reading and just 22 percent in math. Nearly half scored below "Basic" proficiency, struggling with skills as fundamental as using percentages in everyday problems. hese are the students who were eighth graders when COVID-19 closed their schools in 2020. They entered high school amid disruption and are leaving it with fewer academic skills than their predecessors a decade ago.

Meanwhile, a Walton Family Foundation-Gallup study found that between one-quarter and one-half of Gen Z students report that school rarely feels engaging or relevant to their lives. Chronic absenteeism rates remain stubbornly high. learning recovery has stalled, and traditional reforms alone cannot meet the scale of today's challenges.

At the same time, the White House has outlined an ambitious plan for “AI dominance,” focusing on frontier research, advanced chips, and computing infrastructure. Yet the real bottleneck isn’t just compute power—it’s human capital. Today’s students are tomorrow’s workers, engineers, and entrepreneurs. America cannot achieve superintelligence abroad if it is losing basic intelligence at home. The most advanced AI models in the world cannot compensate for a generation unprepared to use them, guide them, or innovate with them.

A generation at risk cannot be the foundation of a nation that hopes to lead in AI. Without a renewed urgency to improve academic performance, the promise of AI leadership will collide with a workforce unequipped to sustain it. America’s AI strategy must be matched by an equally ambitious human capital strategy—one that ensures every student is equipped to learn, adapt, and contribute in an AI-driven world.

The urgency of this moment demands experimentation and innovation - pilots that explore how to harness AI thoughtfully and purposefully in ways that address these challenges. We need more experimentation with AI tutors to personalize instruction and close gaps. We need more tools that utilize generative AI to alleviate the administrative burden that prevents teachers from focusing on their most important work: building deep, meaningful relationships with students, which are the foundation of academic success.

With sustained federal leadership, bold public-private partnerships, and a commitment to responsible design, AI can become one of the most powerful forces for improving learning, supporting teachers, and expanding opportunity. The United States has long led the world in technological innovation; now it must lead in aligning that innovation with human flourishing.