Written testimony submitted to the United States Senate Committee on Health, Education, Labor, and Pensions (HELP) for the Subcommittee on Education and the American Family

June 16th, 2026
Dirksen 430

Hearing on the Future of K12 Education in the Age of Artificial Intelligence

Chairman Tuberville, Ranking Member Blunt Rochester, and distinguished members of the HELP Committee, thank you for convening this critical hearing and for initiating this timely conversation on the rapid integration of artificial intelligence within our nation's K-12 education system.

My name is Erin Mote, and I am the CEO and Co-Founder of InnovateEDU, a non-profit organization dedicated to accelerating innovation in education, improving data interoperability, and ensuring that technology serves as a powerful tool for access and opportunity for all students. I am also the founder of Project Unicorn, a national initiative representing over 1,300 school districts committed to data interoperability and privacy, and the leader of the EDSAFE AI Alliance, a global initiative driving the safe, accountable, fair, and efficacious integration of artificial intelligence in education.

I come before this committee not only as a non-profit executive and a charter school founder, but also as a technologist who has spent more than two decades working at the intersection of technology policy and systemic change. My career has been defined by designing and implementing complex digital infrastructure that is resilient, scalable, and safe. Yet, perhaps my most important perspective is a deeply personal one: I am the mother of school-aged children. Every day, I see firsthand how these technologies impact real kids at the kitchen table and in classrooms, which fuels my urgency to ensure our national digital infrastructure is secure, safe, and intentionally built to serve the next generation.

The technology landscape is undergoing a rapid, structural reorganization. Across our school systems, technology is no longer a static instructional tool; it has become a permanent, foundational layer of infrastructure that shapes how our children and educators work, learn, live, and relate to one another. Data from the Center for Democracy and Technology reveal that approximately 85% of teachers and 86% of students are actively utilizing generative artificial intelligence in their schoolwork.[1] This transition has occurred with unprecedented velocity.

[1] Laird, E., Dwyer, M., & Quay-de la Vallee, H. (2025). Hand in hand: Schools' embrace of AI connected to increased risks to students. Center for Democracy & Technology.

Within a single year, the proportion of school districts that have defined their strategic approach to AI based on specific use cases grew from 31% to 49%. Today, 79% of K-12 school districts report having formal AI guidelines in place, up from 57% the prior year, according to the Consortium for School Networking.[2]

To navigate this landscape, we must reject a pervasive, false dichotomy. I do not believe that safety is the counterpolarity to innovation. Safety is not a drag on progress, nor is it a hurdle we clear at the expense of transformation. Rather, safety is a necessary condition for the use of AI in education—the very guardrail that makes meaningful, sustainable innovation possible. Without it, we mistake rapid tool adoption for progress.

Today, we are seeing the stark consequences of that exact miscalculation play out in real time. Despite this widespread adoption, school districts, educators, and families are experiencing far more than just severe decision fatigue and operational strain. They are confronting the immediate consequences of exposing students to untested technologies—including documented harm to learning, social engagement, and mental health. This crisis has emerged because generative AI entered our classrooms not as an adoption technology that follows a linear lifecycle of institutional procurement, staged pilot programs, and centralized risk evaluations, but as an arrival technology. By bypassing institutional gatekeepers entirely, it has penetrated the educational environment from the outside in through personal consumer habits and direct adoption by students and educators, introducing cognitive and developmental risks before schools can even begin to evaluate them.

According to a national survey by Carnegie Learning, local educational agencies are struggling to keep pace, with 57% of schools and districts failing to provide any professional development on the safe use or monitoring of these tools.[3] Furthermore, local technology leaders express significant anxiety about unfunded state mandates; edtech professionals caution that state-level requirements issued without dedicated funding severely constrain local budgets, leaving schools vulnerable and under-supported. Urgent federal intervention, structured programmatic support, and dedicated funding are required to help school communities, families, and students navigate this transition safely and effectively.

To establish safe and effective classrooms, policymakers and educators must distinguish

---

https://cdt.org/insights/hand-in-hand-schools-embrace-of-ai-connected-to-increased-risks-to-students/

[2] Consortium for School Networking. (2026). *2026 State of EdTech district leadership*. CoSN. https://www.cosn.org/EdTech-topics/state-of-edtech-leadership/

[3] Carnegie Learning. (2025). *The state of AI in education 2025: Key findings from a national survey* [White paper]. https://discover.carnegielearning.com/hubfs/PDFs/Whitepaper%20and%20Guide%20PDFs/2025-AI-in-Ed-Report.pdf

between general-use consumer technologies and purpose-built educational tools. A dangerous blurring of the lines has occurred as school systems rush to integrate consumer-grade AI tools simply because they are cheap (or even free), widely available, and highly publicized. Crucially, the vast majority of software marketed and sold as ed tech today is not actually purpose-built for education; rather, it is merely consumer technology repackaged with an educational layer. However, good enough consumer technology is fundamentally inadequate for our schools. At its core, a consumer-grade large language model (LLM) is optimized for user satisfaction, immediate helpfulness, and platform retention. To maximize commercial engagement, consumer-grade models are engineered to minimize friction, thereby defaulting to a dangerous form of algorithmic sycophancy. They do not teach; they indulge. By routinely validating incorrect premises and directly supplying unearned answers, they prioritize relentless platform retention over the productive struggle required for true cognitive development.

By contrast, learning is an active, constructivist process that requires productive struggle. True purpose-built educational technology (EdTech) must be designed to scaffold learning, adjusting tasks to a student's Zone of Proximal Development (ZPD) to build independent critical thinking and conceptual retention.[4] Examples of this type of purpose-built educational technology are being implemented in Chairman Tuberville's home state, where districts across Alabama are utilizing AI-driven math accelerators, such as Alabama Math, Science, and Technology Initiative (AMSTI)- aligned tools. These systems dynamically adjust the complexity of multi-step word problems based on real-time student performance, helping teachers rapidly identify specific foundational gaps in early algebra.

Research increasingly demonstrates that artificial intelligence is most effective in education not as a standalone replacement, but as a catalyst that amplifies a human-centered learning experience. By offloading administrative functions (such as data diagnostics) to technology, human educators are freed to focus on what they do best: emotional support, motivation, and scaffolding. A comprehensive review by Stanford University supports this, finding that AI tools that provide automated feedback directly to human tutors significantly improved both instructional quality and student learning outcomes. However, Stanford researchers also cautioned that without human-guided instruction, fully autonomous AI can reduce a student's cognitive burden and hinder long-term independent mastery.[5] Complementing this, the educational platform Eedi leverages AI to analyze student error patterns, putting these targeted diagnostic insights directly into teachers' hands and increasing classroom efficacy by 27%.

---

[4] Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.

[5] Fesler, L. (2026). *The Evidence Base on AI in K-12: A 2026 Review*. Stanford SCALE Initiative. https://scale.stanford.edu/sites/default/files/The%20Evidence%20Base%20on%20AI%20in%20K-12%20Report.pdf

Ultimately, combining AI's analytical insights with the irreplaceable empathy and intuition of a human mentor creates conditions for deep, durable learning.[6]

Conversely, when consumer-grade systems are introduced as classroom substitutes, they contribute to cognitive surrender, a state where students outsource high-level synthesis and conceptual analysis to algorithms.[7] Stanford's SCALE Initiative documented a clear performance cliff resulting from this outsourcing: while students using unrestricted AI show immediate performance improvements during active tasks, their independent performance plummets once the tool is removed.[8] As this committee evaluates the integration of generative artificial intelligence, we must confront a dangerous, self-reinforcing psychological loop. At Arizona State University, Dr. Punya Mishra has identified what he calls The Yes-Bot Problem, which is the inherent tendency of generative AI to be unceasingly agreeable, acting like an over-eager, sycophantic intern that confirms a user's ideas rather than challenging them.[9] When this agreeable technology meets human psychology, it triggers what Wharton researchers call cognitive surrender. Because the AI functions as a fluent, external cognitive layer, users completely relinquish their independent reasoning. The empirical data is alarming: when an AI chatbot intentionally gave incorrect answers, users accepted the error 80% of the time.[10] Worse, their internal confidence actually *increased* by nearly twelve percent because they mistook the Yes-Bot's authoritative, accommodating tone for actual correctness. We are actively subsidizing a dangerous psychological loop: we are placing Yes-Bots engineered to tell users exactly what they want to hear in front of neurologically vulnerable youth, conditioning a generation to surrender the critical discernment required to question them.

When we allow consumer-grade AI and technology in our classrooms, we introduce significant educational, cognitive, and privacy risks. Purpose-built EdTech does the exact opposite: it fosters active learning by prioritizing measurable growth, utilizing Socratic methods to prompt reflection, and anchoring progress to actual skill mastery. The risk extends from the cognitive to the structural. While purpose-built educational technology operates under formal institutional contracts that ensure strict compliance with FERPA and COPPA, consumer platforms subject our

---

[6] LearnLM Team Google, Wang, A., Rysbek, A., Huber, A., Nambiar, A., Kenolty, A., Caulfield, B., Lilley-Draper, B., Groot, B., Veprek, B., Burdett, C., Willis, C., Barton, C., Smith, D., Mu, G., Walters, H., Jurenka, I., Hulls, I., Stalley-Moores, J., . . . Brazão, V. (2025). AI tutoring can safely and effectively support students: An exploratory RCT in UK classrooms. *arXiv*. https://doi.org/10.48550/arXiv.2512.23633

[7] Shaw, S. D., & Nave, G. (2026, January 11). Thinking-Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender. *The Wharton School Research Paper*. https://doi.org/10.31234/osf.io/yk25n_v1

[8] Fesler, L., Martinez Claeys, J., Agnew, C., & Loeb, S. (2026, March 23). *The evidence base on AI in K-12: A 2026 review*. Stanford SCALE Initiative.

[9] Mishra, P. (2025, October 21). *The yes-bot problem: Why agreeable AI makes learning harder*. punyamishra.com/2025/10/21/the-yes-bot-problem-why-agreeable-ai-makes-learning-harder/

[10] See Note 5.

children to broad public terms that allow their data to be scraped for training future models. We must establish a clear boundary that protects student privacy and preserves cognitive development from the incentives of commercial tech.

To resolve the systemic mismatch between commercial offerings and classroom needs, school districts must transition from volume-based procurement toward evidence-based purchasing. This transformation is structured around the Five EdTech Quality Indicators established by a coalition of student, family, and educator-centric organizations. When school systems and states anchor their procurement standards to these indicators (Safe, Interoperable, Usable, Inclusive, and Evidence-Based), they establish a consistent baseline of pedagogical and ethical integrity.

By providing this unified framework, the coalition offers a much-needed antidote to a growing crisis in our education system: we simply have too many apps in America's classrooms, creating a fragmented, unmanageable digital landscape that frequently overwhelms both teachers and students. According to the 2025 EdTech Top 40 Report, school districts accessed an average of 2,982 distinct digital EdTech tools over the school year. Every month, this digital ecosystem remains incredibly dense, with districts seeing students interact with an average of 1,165 tools and teachers interact with 1,045 tools. When looking at individual footprints over the entire academic year, a single student now navigates a core set of 48 unique tools to power their learning, while educators are forced to rely on an average of 50 different apps to manage their instruction, data, and daily classroom workflows.[11]

To cut through the digital noise and move away from chaotic procurement, school systems must anchor their technology decisions in a limited, definitive set of foundational measures, ensuring that tools are systematically vetted as Safe, Interoperable, Usable, Inclusive, and Evidence-Based. At the heart of this necessary curation must be a commitment to the federal ESSA (Every Student Succeeds Act) tiers of evidence. The stark contrast in market readiness underscores the absolute necessity of this gatekeeping. While approximately 40% of purpose-built EdTech tools meet a validated ESSA tier of evidence, only 2% of general consumer tools do. Meeting at least a Tier 4 within the ESSA framework must become the non-negotiable ticket to enter America's classrooms.[12] Having an informed hypothesis about a tool's specific impact on children, instruction, and educator pedagogy must serve as the structural underpinning of our shift toward an evidence-informed, developmentally and pedagogically

---

[11] Instructure. (2025). The EdTech Top 40: K-12 EdTech Engagement During the 2024-25 school year. https://www.instructure.com/edtech-top40Every month

[12] Instructure, & InnovateEDU. (2026). *2026 EdTech evidencethereby forcing a healthy realignment of market incentives and report*. https://drive.google.com/file/u/2/d/1YKtmaiU-Ux-SIIhwSRNvH6_6w4_ag39o/view?usp=drive_link

appropriate use of technology — or a new ROI: the return on instruction.

True return on instruction is impossible to measure without clear, standardized metrics, which is why state leaders are stepping in to build the infrastructure required for such evaluation. In a pioneering move, Ohio is building a statewide EdTech repository powered by the ISTE EdTech Index, a partner with InnovateEDU on the 5 Quality Indicators work, to provide district decision-makers and educators with unprecedented transparency into the digital tools used in classrooms. Grounded in five critical EdTech quality indicators, this centralized platform puts actionable data regarding privacy, data interoperability, and evidence of efficacy directly into the hands of local stakeholders. By streamlining the evaluation process, the initiative empowers educators to make safer, more effective instructional choices while fostering community trust in school technology investments. This proactive approach directly confronts a systemic vulnerability that has plagued schools for years. Currently, the procurement process in K-12 education is broken: districts routinely purchase software licenses based on marketing promises or general popularity. To solve this, school systems must shift toward outcomes-based contracting (OBC). Under an OBC model, a portion of vendor payments is contingent on the technology demonstrating measurable improvements in student outcomes, such as math or reading proficiency.

Outcomes-Based Contracting shifts financial risk away from the American taxpayer and ties vendor compensation directly to empirical student growth, thereby forcing a healthy realignment of market incentives and transforming vendors from software providers into accountable partners in student flourishing. Early pilots, such as the partnership between Duval County Public Schools and an AI tutoring organization, demonstrate that tying payment to measurable student growth changes provider relationships, strengthens implementation, and helps schools sustain interventions beyond short-term funding cycles.[13]

But updating our procurement models is only half the battle; we must also modernize the outdated metrics we use to evaluate classroom technology. For too long, the conversation around technology in K-12 education has been limited by the blunt metric of screen time, which is an outdated focus on counting minutes rather than measuring impact. As we navigate a transformative era of digital infrastructure and artificial intelligence in the classroom, simply monitoring hours is no longer a useful proxy for learning. We must urgently shift our focus to screen value, evaluating whether digital tools drive active engagement, foster deep pedagogical impact, and adhere to safe-by-design principles. The critical question for leaders and educators is not how long a student is looking at a device, but whether that interaction is meaningful,

[13] Noakes, S., Shell, A., & Vollavanh, A. (2026, March). *Instructional Coherence and Shared Responsibilities: OBC's Impact on EdTech Procurement*. *Digital Promise.* https://doi.org/10.51388/20.500.12265/288

instructionally anchored, and empowering. By anchoring our strategy in value rather than time, we can move past passive consumption and design digital ecosystems that truly accelerate student learning and systemic change.

Take literacy as a stark example of why evidence must shape not just *if* we use technology, but exactly *how* and *when* we use it. For building deep reading comprehension, evidence shows that screen-based reading is six to eight times less effective than reading physical books—with the critical exception of students with disabilities, for whom digital text serves as a vital accommodation. The data make it clear that, for teaching reading comprehension, we should buy physical books.[14]

But now consider decoding, an essential foundational skill required before comprehension can even occur. Young readers need a thousand at-bats with decoding to achieve fluency, much like mastering multiplication tables. This is precisely where purpose-built educational technology excels: it is uniquely engineered for this type of engaging, high-repetition, and pedagogically anchored practice. Grounding our deployment in this level of nuance ensures that we use technology as a precise tool to accelerate specific cognitive skills, rather than treating screens as a blanket substitute for authentic learning.[15]

If we want to see what measuring true impact looks like, we need only look at the findings from Duval County. By shifting the focus to high-quality implementation, an initial WestEd study of 2,000 students showed that targeted AI tutoring under an OBC made second-graders 28 percentage points more likely to meet ELA proficiency. These results clearly indicate that purpose-built technology, when deployed correctly, can fundamentally change student trajectories. Yet no effects were found in older grades despite receiving the same dosage. This critical nuance shows that technology is not a silver bullet; its success depends entirely on developmental readiness, meaning our frameworks must prioritize targeted interventions in which purpose-built EdTech actually drives measurable learning gains. And this isn't just in Florida; Louisiana is leveraging adaptive AI software to support statewide foundational literacy initiatives.[16] The software analyzes individual phonetic mastery and tailors micro-lessons for students who need extra reinforcement, aligned with the state's rigorous early reading frameworks.

---

[14] Delgado et al.. (2018, November). *Don't throw away your printed books: A meta-analysis on the effects of reading media on reading comprehension.* Educational Research Review, 25, 23-38. https://doi.org/10.1016/j.edurev.2018.09.003

[15] Silverman, R. D., Keane, K., Darling-Hammond, E., & Khanna, S. (2025). *The Effects of Educational Technology Interventions on Literacy in Elementary School: A Meta-Analysis*. Review of Educational Research. https://doi.org/10.3102/00346543241261073

[16] Vesco, B. (2024, July 25). *Louisiana to roll out AI reading tutor in 25 school districts*. Government Technology. https://www.govtech.com/education/k-12/louisiana-to-roll-out-ai-reading-tutor-in-25-school-districts

Active, collaborative, and Socratic interactions grounded in instructional design are highly beneficial. In contrast, passive, consumer-focused scrolling designed to maximize attention is linked to harm to cognitive and social development. This crucial distinction between passive consumption and active engagement forms the foundation for addressing the next major shift in digital technology. Yet, our current public conversation often ignores this nuance entirely. Focusing on screen time rather than screen value is a blunt, short-sighted approach that ultimately harms children, stalls educational progress, and stifles learning. By treating all digital interactions as inherently toxic, strict limits fail to distinguish between passive, mindless consumption and active, high-value educational engagement. Moreover, this rigid focus disproportionately harms the one in five children in America with learning disabilities and attention issues, for whom educational technology is not a luxury, but often a critical lever for mandated accommodations.[17] When states and districts implement sweeping, restrictive policies based entirely on clock hours, they inadvertently strip away essential digital infrastructure such as text-to-speech, adaptive screen readers, and personalized pacing tools that level the playing field. Ultimately, these blunt instruments cut students off from vital learning frameworks, leaving them less prepared for a technology-driven world and widening learning gaps for the very learners who rely on tailored digital tools to thrive.

True protection, therefore, cannot be achieved through a strategy of digital avoidance or categorical bans; instead, we must pivot toward proactively preparing our school communities to engage with these emerging tools through a sharp, discerning lens. As artificial intelligence increasingly shapes our daily lives, fostering AI literacy and establishing a mindset of critical consumerism among students, educators, and families is essential for navigating the digital world responsibly. True literacy goes beyond knowing how to use a tool; it requires an understanding of data privacy, algorithmic bias, and real-world impacts, so that users become thoughtful evaluators rather than passive consumers of technology.

Forward-thinking school districts and states are already shifting their curricula to build these vital competencies. In Ranking Member Blunt Rochester's home state, for example, the Delaware Department of Education, in partnership with an EDSAFE AI Alliance state policy lab, has published a guide that connects traditional media literacy standards with generative artificial intelligence to help educators teach students to navigate AI-driven content and tools. The state's work helps ensure that Delaware students develop critical thinking skills, ethical online behaviors, and the ability to detect algorithmic bias or misinformation. In another of the

---

[17] Chung, A.-M., & Sharma, K. (2026, April 10). *Blanket EdTech bans miss the mark: Better policies for student tech use*. New America.
https://www.newamerica.org/insights/blanket-edtech-bans-miss-the-mark/

EDSAFE AI Alliance policy labs, the Lake Washington School District, AP Computer Science Principles and AP Computer Science A courses are adding new units on artificial intelligence, machine learning, cybersecurity, and other emerging technologies to modernize their curricula. Through these additions, students gain relevant, up-to-date technical skills and are fully equipped to succeed in future technology careers. Cultivating this critical lens also means introducing these concepts early in a student's academic journey; for instance, Juanita High School is introducing an AI Foundations course for the next school year, allowing students to explore artificial intelligence concepts, ethical considerations, and real-world applications. By blending technical proficiency with ethical scrutiny, educators ensure that the next generation is prepared not only to build the tools of tomorrow but also to question and shape their place in society critically.

Achieving this level of critical agency, however, requires moving past basic applications of artificial intelligence. While early EdTech discussions often limit AI's role to automated tutoring or step-by-step homework help, the technology's most transformative uses emerge when it supports open-ended, self-directed learning models like project-based inquiry. A real-world example of this is happening in Vermont, where districts across the state are leveraging the Vermont Agency of Education's statewide AI framework (issued in early 2026). Because of its long-standing legislative commitment to Personalized Learning Plans (PLPs) and proficiency-based graduation requirements, the state's official guidance encourages high schools to treat generative AI as an active "co-creator" and cognitive "thinking partner" rather than an intellectual shortcut.[18]

A standout example of this positive integration in action is the Essex Westford School District (EWSD) in Vermont, which proactively hired specialized digital learning leaders to anchor this macro framework in daily practice. Through collaborative educator workshops, the district brings teachers onto the same page to determine how to integrate AI into their teaching and assessment strategies appropriately. Simultaneously, EWSD builds critical AI literacy among students, equipping them to understand how these systems operate while keeping human values, ethics, and critical judgment at the center. By ensuring students have access to vetted, safe digital learning technologies, the district empowers them to innovate safely and truly drive their own learning by using AI as a structured research partner to brainstorm open-ended concepts, locate community resources, and organize self-directed capstone projects that align with local sustainability and civic goals.

Yet, this vision of AI as a structured, pedagogy-first learning partner stands in alarming contrast to an unvetted consumer marketplace where technology ceases to serve as an instructional tool

[18] Vermont Agency of Education. (2026). Vermont AI guidance for education: A practical framework to support local decision-making. State of Vermont. vermont.gov

and instead begins to simulate human intimacy. When a technology introduces urgent and foreseeable harm to young people, we are systematically and morally obligated to intervene, and this is precisely the crisis we face with the rise of AI companions and chatbots that simulate human relationships. For developing minds navigating identity and belonging, these platforms present a profound psychological risk by mimicking empathy, friendship, and unconditional validation.[19] Authentic human relationships demand the difficult work of navigating conflict, reciprocity, and mutual vulnerability. An AI companion offers none of this. It exploits the empathy gap, delivering a frictionless, synthetic intimacy that accelerates adolescent social isolation and distorts a developing child's emotional baseline.[20] When technology shifts from an interactive tool to a surrogate for human connection, we cannot afford a passive wait-and-see approach; establishing strict safe-by-design guardrails and critical AI literacy is an immediate imperative to protect youth cognitive and social development.

The introduction of unstructured AI companions into minors' lives poses unique psychological and developmental risks. Because adolescents are neurologically vulnerable due to a developing prefrontal cortex, they are highly susceptible to anthropomorphizing conversational agents. They regularly mistake programmed empathy for genuine human connection, risking unhealthy attachments. A July 2025 Linewize survey revealed that 60% of U.S. schools reported students confiding in AI chatbots rather than teachers or parents, and 45% reported students forming emotional attachments to AI companions. These frictionless, always-agreeable digital relationships lack the real-world compromises and conflict resolution necessary to build genuine social-emotional skills. In the most tragic instances, such as the widely documented suicide of 16-year-old Adam Raine, unregulated consumer AI models have validated self-harm ideation, discouraged users from seeking human help, and directly contributed to psychological crises. This empathy gap highlights why guardrails must be established: a chatbot's expression of emotion is a programmed prediction, not a felt experience.[21]

Securing this vital balance between protection and privacy is a technical necessity, but it cannot be architected in a socioeconomic vacuum. If our safety frameworks create overly complex, high-friction compliance hurdles, they risk placing essential innovations completely out of reach for under-resourced communities. Technological policies must be intentionally designed to avoid inadvertently exacerbating the digital divide. For students with disabilities, adaptive AI technologies, including speech synthesis for speech impairments, specialized conversational

---

[19] Namvarpour, M., Brofsky, B., Medina, J., Akter, M., & Razi, A. (2025). Understanding teen overreliance on AI companion chatbots through self-reported Reddit narratives. *arXiv*. https://doi.org/10.48550/arXiv.2507.15783

[20] Kurian, N. (2023). AI's empathy gap: The risks of conversational Artificial Intelligence for young children's well-being and key ethical considerations for early childhood education and care. Contemporary Issues in Early Childhood, 26, Article 1. https://doi.org/10.1177/14639491231206004

[21] Linewize. (2025, December 10). *See the signs: What U.S. school districts are seeing and doing about student mental health in the digital age*. Qoria.

chatbots for autism spectrum disorder, and annotated text for dyslexia, serve as critical portals to basic educational opportunity. This same principle is opening doors in Senator Kim's home state. In New Jersey, schools are using AI platforms to support multilingual learners. The tools provide real-time translation and contextual vocabulary scaffolding across complex science and social studies texts, allowing English language learners to master rigorous grade-level content while simultaneously building language proficiency.

However, the promise of these tools is deeply constrained by ongoing infrastructure gaps. A landmark analysis by All4Ed, UnidosUS, the National Indian Education Association and the National Urban League, using U.S. Census Bureau American Community Survey data, found that roughly one in three Black and Latino households and more than 36% of rural households lack high-speed internet or a dedicated personal computer.[22] This access gap forces these students to rely solely on school-issued devices and networks for their education.

One state, Kentucky, through its partnership with the EDSAFE AI Alliance policy labs, is working to bridge these infrastructure gaps proactively. The Kentucky Educational Development Corporation (KEDC) has launched the AI Rolling Lab, a mobile classroom designed to bring hands-on artificial intelligence learning experiences directly to students and educators across the state. This ensures that all districts, regardless of their size or location, receive access to cutting-edge technology and future-focused career preparation.

To protect these students' civil rights, federal safety policies must act as a scalpel, not a sledgehammer. Blunt, heavy-handed mandates will inevitably create compliance bottlenecks that cut off access for rural and low-income students who depend entirely on school-provided tech. Congress can prevent this by adopting a waterfall approach to federal safety frameworks. The waterfall approach applies lower-friction, privacy-preserving standards for general learning environments and purpose-driven instructional models. Rather than a blanket, one-size-fits-all mandate, this strategy applies heavy-handed safety and privacy controls to high-risk applications at the top, such as open-ended public consumer platforms or chatbots that simulate human relationships, while letting the regulatory burden drop for purpose-driven educational tools. For lower-risk environments, such as a specialized speech synthesis tool or a dyslexia text annotator, the framework uses streamlined, privacy-preserving standards that protect student data without creating bureaucratic roadblocks. This distinction is critical for access and opportunity; if a single, blunt law treats a classroom accessibility tool with the same severity as an unregulated consumer chatbot, sky-high compliance costs will disproportionately harm under-resourced rural and urban school districts that lack the massive IT and legal teams to clear those hurdles. Ultimately, a waterfall approach ensures that safety guardrails

[22] Alliance for Excellence Education, National Urban League, National Indian Education Association, & UnidosUS. (2020, August). *Students of Color Caught in the Homework Gap*. All4Ed. https://futureready.org/wp-content/uploads/2020/08/HomeworkGap_FINAL8.06.2020.pdf

aggressively target genuine harms at the top while allowing essential, day-to-day learning tools to flow smoothly to the students who need them most. This tiered regulatory strategy ensures that well-intentioned compliance mandates do not inadvertently lock rural students in Texas, Alabama, and across the nation out of modern digital resources, keeping the doors of opportunity wide open.

I believe that to navigate these challenges, the federal government must enact immediate, structural policy changes to protect children, support families, empower educators, and maintain our global competitiveness.

The federal government must reestablish a coordinated center of gravity for educational technology policy. Congress should reconstitute and fully staff the Office of Educational Technology (OET) within the U.S. Department of Education, or create a similarly situated office at the Department of Health and Human Services (HHS) or the Department of Labor. This central office would lead national efforts to evaluate classroom technology, publish plain-language safety guides for families, and manage a National Child-AI Safety in Schools Lab. By building this public infrastructure, the federal government can lower barriers for small startups while shielding school systems from unvetted consumer products. Crucially, this office must work in concert with the Department of Commerce NIST function and the newly constituted Center for AI Standards and Innovation (CAISI) to establish standard evaluation metrics, ensuring that public investments are directed exclusively toward technologies verified for safety, accessibility, and pedagogical efficacy.

Establishing these baseline metrics is an essential shield. Yet the uncomfortable reality is that this technology has already become widespread in America's classrooms long before we have definitive answers. The velocity of AI deployment has far outpaced the scientific community's understanding of its long-term developmental effects. While AI tools flood classrooms, we are operating in an information vacuum: According to Stanford's SCALE initiative, there are zero high-quality causal studies on student AI use in U.S. K-12 schools, leaving us with short-term data that completely ignores long-term learning, equity, and student wellness.[23] Congress must expand funding and direct the National Science Foundation (NSF), the Institute of Education Sciences (IES), and the National Institutes of Health (NIH) to establish a joint, interdisciplinary research agenda.

This federal initiative must focus on:

- The cognitive and neural consequences of sycophant behavior and immediate algorithmic feedback on student memory consolidation and retention.
- The psychological and developmental impacts of long-term adolescent engagement with

[23] See Note 8.

AI companion chatbots.
- Creating the empirical evidence base required to define and measure screen value as opposed to simple screen time.
- Evaluating the long-term impact of automated, stealth assessments on student academic performance and emotional well-being.

A national research agenda must prioritize child safeguarding, focusing on preventing algorithmic exploitation, digital deception, and the generation of non-consensual imagery. While the federal government has historically struggled to regulate online domains, the bipartisan TAKE IT DOWN Act provides a vital blueprint by establishing a federal notice-and-removal standard for non-consensual AI-generated imagery and digital forgeries. Enforcing these swift removal standards will require substantial federal funding, yet the ultimate success of child safeguarding hinges on implementing proactive, privacy-respecting technical solutions. Furthermore, the national research agenda should define technical standards for Zero-Knowledge Proofs (ZKPs) and Tokenized Trust models, enabling robust, privacy-preserving age assurance (PPAA) that protects children from adult-content platforms without subjecting them to intrusive biometric surveillance. Privacy-preserving age verification is essential because it strikes the vital balance between safety and autonomy, allowing platforms to robustly protect children from online harms without forcing users to trade away their digital privacy.

Safeguarding individual privacy at the user level, however, is futile if the broader institutional networks housing this data are left entirely vulnerable to hostile actors. Cybersecurity remains the absolute priority for school technology leaders, according to the COSN State of EdTech District Leadership Survey. School districts are increasingly targeted by sophisticated, AI-enabled cyberattacks, with over a quarter of districts classifying network and administrative disruptions as severe, leaving systems completely offline for more than 24 hours. Paradoxically, 65% of districts operate without a dedicated cybersecurity budget, and 51% face escalating insurance premiums that drain funds from technical staffing.[24] Congress must aggressively protect the Federal Communications Commission's (FCC) E-Rate program, which helps schools secure basic connectivity, filtering, and content monitoring. Additionally, the federal government should expand the E-Rate cybersecurity pilot program to provide dedicated, predictable funding for network defense, endpoint protection, and advanced firewalls, preventing the siphoning of classroom funds to address security emergencies.[25]

According to research by the Center for Democracy & Technology and national educator surveys, teachers play an outsized role in managing technology: 45% receive automated activity alerts weekly, 17% receive them daily, and 65% are directly responsible for following up on

[24] Consortium for School Networking. (2023). *CoSN's annual K-12 IT leadership survey report*.
[25] Consortium for School Networking. (2026). *2025 State of EdTech district leadership*.

student flags.[26] Despite these immense obligations, only 31% of teachers have received formal guidance on managing monitored student data privately and securely, and 69% report a lack of substantial training in classroom AI integration.[27] To bridge this vacuum, Congress must fully fund and expand the allowable uses of Title II, Part A (Supporting Effective Instruction) and Title IV, Part A (Student Support and Academic Enrichment) formula grants.

- Train educators in the cognitive science of technology integration to help them identify the empathy gap and prevent cognitive surrender.
- Compensate educators for the hidden labor of auditing AI outputs, participating in school-level tool evaluations, and managing academic integrity.
- Provide comprehensive professional development that equips teachers to become critical consumers of classroom tools.

To operationalize these recommendations, the federal government must support innovative state-level pilots that move beyond uncoordinated, experimental software adoption and align with benchmarks firmly anchored in the learning sciences. Exemplifying this shift, Impact Florida, in partnership with the University of Florida, has launched SPARK, an 18-month initiative to help FL school districts thoughtfully, responsibly, and effectively integrate artificial intelligence into K-12 classrooms through work with the EDSAFE AI Alliance. This programmatic focus shifts districts from a reactive mindset to an intentional path, allowing them to establish cohesive AI policies, explore outcomes-based contracting models, and scale tested instructional practices for long-term success. Echoing Florida's work, Utah's Educational Technology Regulatory Sandbox offers a powerful real-world model for this structural shift. Designed to move away from reactive tech procurement, Utah's framework creates a supervised, time-limited testing environment that mandates strict evidence-based evaluations under real classroom conditions before any tool can be scaled statewide.[28]

This proactive state-level gatekeeping responds directly to urgent warnings from the Brookings Global Task Force on AI in Education. In their report *A New Direction for Students in an AI World*, Rebecca Winthrop and her colleagues caution against AI-diminished learning, which is a state where an overreliance on automated tools short-circuits the productive struggle students need to develop critical thinking and their own unique voices.[29] Strikingly, this research reveals

---

[26] Laird, E., Grant-Chapman, H., Venzke, C., & Quay-de la Vapeople'sllee, H. (2022). *Hidden harms: The misleading promise of monitoring students online*. Center for Democracy & Technology. https://cdt.org/insights/report-hidden-harms-the-misleading-promise-of-monitoring-students-online/

[27] Youngstown State University. (2024). National educator AI readiness and confidence survey. School House / YSU Research Initiative.

[28] Burns, M., Winthrop, R., Luther, N., Venetis, E., & Karim, R. (2026). A new direction for students in an AI world: Prosper, prepare, protect. Brookings Institution. https://www.brookings.edu/articles/a-new-direction-for-studenFlorida'sts-in-an-ai-world-prosper-prepare-protect

[29] See Note 8.

that young people’s primary concern regarding AI isn't future job displacement, but cognitive loss, which is the literal fear of losing their own capacity to think deeply by coasting in a passive passenger mode.

To prevent this, pilots must ground their evaluations in frameworks such as the Learning Sciences Benchmark, being developed by the EDSAFE AI Alliance in collaboration with Arizona State University. Unlike standard industry metrics that evaluate an AI model purely on processing speed or technical accuracy, this framework establishes a pedagogical yardstick. By measuring tools based on their direct impact on human cognition and student agency, this standardized mechanism ensures that classroom AI systems serve as intentional scaffolds that cultivate an active-explorer mindset, rather than as simple automation tools that accelerate the very cognitive offloading that students themselves fear.

On a national scale, the Global EdTech Trialing Network (GETN-US), supported by InnovateEDU, has synthesized these models to advocate for a federal AI safety and efficacy testbed infrastructure. Collaborative safety and efficacy trial environments, like those pioneered in New York City Public Schools, allow educators and school districts to test implementation alongside technology. They provide the public infrastructure we need to verify EdTech efficacy, protecting public budgets and securing America's cognitive and economic future.

The rapid integration of generative AI into American classrooms is not a temporary technological trend; it is a fundamental re-architecting of how we develop future generations to be productive, independent, and engaged citizens. We face a definitive, generational choice. We can passively surrender our classrooms to the cheapest, most sycophantic algorithms on the market, outsourcing our children's critical thinking, executive functioning, and social-emotional development to commercial platforms.

Or, we can seize this historic inflection point to declare that technology must serve humanity, acting as a powerful partner in human flourishing rather than a replacement for it. By fully funding our schools' digital defense, teacher training, and research infrastructure, we protect the sacred spark of human potential and ensure a secure and prosperous future for every American child.