

# A NEW DIRECTION FOR STUDENTS IN AN AI WORLD: PROSPER, PREPARE, PROTECT

MARY BURNS, REBECCA WINTHROP, NATASHA LUTHER,  
EMMA VENETIS AND RIDA KARIM



# A NEW DIRECTION FOR STUDENTS IN AN AI WORLD: PROSPER, PREPARE, PROTECT

JANUARY 2026

**Mary Burns** is an independent research consultant for the Center for Universal Education, Global Economy and Development program, Brookings Institution

**Rebecca Winthrop** is a Senior Fellow and the Director at the Center for Universal Education, Global Economy and Development program, Brookings Institution

**Natasha Luther** is a Senior Research Analyst at the Center for Universal Education, Global Economy and Development program, Brookings Institution

**Emma Venetis** is a Senior Research Analyst at the Center for Universal Education, Global Economy and Development program, Brookings Institution

**Rida Karim** is a college student and Communications Coordinator for the National Student Board Member Association

## **ACKNOWLEDGEMENTS**

We would like to thank all our steering committee members, resource experts, study participants, and others who contributed to this research. We are deeply grateful for their time, insights, and generosity and without them this report would not have been possible. Please see Annex B for our full acknowledgements.

We are grateful to Ashleigh Ekwenugo, Mashhood Bhat, and Izzy Taylor for their contributions to this report. We are also grateful to our reviewers including Emily Morris, Punya Mishra, and Jennifer O'Donoghue. Design was done by Marian Licheri Hougaard and Andreina González Anzola.

The Brookings Institution is a nonprofit organization devoted to independent research and policy solutions. Its mission is to conduct high-quality, independent research and based on that research, to provide innovative, practical recommendations for policymakers and the public. The conclusions and recommendations of any Brookings publication are solely those of its author(s), and do not reflect the views or policies of the Institution, its management, its other scholars, or the funders acknowledged below.

Brookings gratefully acknowledges the support of EY and the Lego Foundation.

## **ABOUT THE CENTER FOR UNIVERSAL EDUCATION**

The Center for Universal Education (CUE) is a leading policy center focused on universal quality education and skills development around the world. CUE helps inform the development of policy related to global education and promotes actionable strategies for governments, civil society, and private enterprise.

Throughout the drafting of this report, the authors used generative AI tools such as Anthropic's Claude, Microsoft's Co-pilot, and OpenAI's ChatGPT. These tools were used, in limited cases, to summarize research articles, locate sources, and condense and edit text. The outputs were reviewed and revised for factual accuracy and checked for plagiarism by the authors prior to publication. The AI tools did not contribute any original ideas to the report.

# Table of Contents

- LIST OF ACRONYMS ..... 6
- FIGURES ..... 7
- BOXES ..... 8
- EXECUTIVE SUMMARY ..... 10
- I INTRODUCTION ..... 16
  - This report: Its purpose and main message ..... 17
  - Foundational premises of this report ..... 21
  - Our methods: A premortem on AI and children’s education ..... 22
  - What is generative AI? ..... 24
- II CONTEXT: Students are increasingly accessing AI outside of school, and families continue to be on the front lines ..... 27
  - Widespread AI use by children and youth ..... 28
  - The blurred lines of children’s AI use ..... 29
  - Families on the front lines ..... 31
- III POTENTIAL OUTCOME 1: AI can enrich learning if well-designed and anchored in sound pedagogy ..... 33
  - Benefit 1: AI can improve equity by addressing educational resource gaps and expanding access to education ..... 34
  - Benefit 2: AI can optimize teacher time for greater focus on students ..... 38
  - Benefit 3: AI can improve student learning ..... 41
  - Benefit 4: AI can tailor learning to each student’s needs ..... 44
  - Benefit 5: AI can extend learning to neurodivergent students and students with disabilities ..... 47
  - Benefit 6: AI can advance assessment ..... 50
- IV POTENTIAL OUTCOME 2: AI can diminish learning if overused with no guardrails ..... 53
  - Risk 1: AI can undermine students’ cognitive development ..... 55
  - Risk 2: AI can impede students’ social and emotional development ..... 69
  - Risk 3: AI can degrade trust in education ..... 83
  - Risk 4: AI can threaten students’ safety ..... 88
  - Risk 5: AI dependence can erode students’ autonomy and agency ..... 104
  - Risk 6: AI can deepen equity divides ..... 110

V

<b>RISKS VERSUS BENEFITS: Currently, AI's risks overshadow its benefits</b> . . . . .	<b>119</b>
AI's benefits are numerous . . . . .	119
But today, these benefits are overshadowed by risks . . . . .	120
Advancing AI learning experiences that enrich, not diminish . . . . .	120
Drivers of AI risks . . . . .	123
The time to act is now . . . . .	123

VI

<b>IT'S NOT TOO LATE: Recommendations for mitigating risks and harnessing benefits</b> . . . .	<b>124</b>
We all have a role to play . . . . .	127
Prosper . . . . .	128
Prepare . . . . .	142
Protect . . . . .	151
Conclusion: Navigating the crossroads . . . . .	163

<b>ENDNOTES</b> . . . . .	<b>165</b>
---------------------------	------------

<b>REFERENCES</b> . . . . .	<b>166</b>
-----------------------------	------------

<b>ANNEX A: Example AI literacy frameworks</b> . . . . .	<b>203</b>
----------------------------------------------------------	------------

<b>ANNEX B: Acknowledgements</b> . . . . .	<b>210</b>
--------------------------------------------	------------

<b>ANNEX C: Research methods</b> . . . . .	<b>216</b>
--------------------------------------------	------------

# LIST OF ACRONYMS

AAC	Augmentative and Alternative Communication
AI	Artificial Intelligence
API	Application Programming Interface
AR	Augmented Reality
ASCD	Association for Supervision and Curriculum Development
ASD	Autism Spectrum Disorder
AUP	Acceptable Use Policy
AWE	Automated Written Evaluation
BERT	Bidirectional Encoder Representations from Transformers
BOT	Build-Operate-Transfer model
CDT	Center for Democracy & Technology
CFF	Cognitive Forcing Functions
COPPA	Children's Online Privacy Protection Act
CPRA	California Privacy Rights Act
FERPA	Family Educational Rights and Privacy Act
GDC	Global Digital Compact
GDPR	General Data Protection Regulation
GPT	Generative Pre-Trained Transformer
IFE	Innovative Financing for Education
ISTE	The International Society for Technology in Education
ITS	Intelligent Tutoring System
L2	Second Language
LARF	Let AI Read First
LLM	Large Language Model
MAiLS	Machine Agents in the Learning Scale
MIT	Massachusetts Institute of Technology
MMLU	Massive Multitask Language Understanding
NOLAI	Netherlands National Education Lab AI
OCR	Optical Character Recognition
PPP	Public-Private Partnership
RBD	Regulation by Design
RBF	Results-Based Financing
SLM	Small Language Model
STEM	Science, Technology, Engineering, and Math
UK	United Kingdom
UNESCO	United Nations Educational, Scientific, and Cultural Organization
U.S.	United States
VR	Virtual Reality
WEIRD	Western, English-speaking, Industrialized, Rich, and Democratic
XAI	Explainable Artificial Intelligence

# FIGURES

- FIGURE 01: Potential benefits and risks of AI in education. . . . . 19
- FIGURE 02: Bending the arc from diminished to enriched AI learning experiences . . . . . 20
- FIGURE 03: Map of study participants' countries . . . . . 24
- FIGURE 04: Children's AI use around the world . . . . . 29
- FIGURE 05: The blurred lines of children's AI use . . . . . 30
- FIGURE 06: ChatGPT usage by month . . . . . 31
- FIGURE 07: Risks identified by study participants . . . . . 54
- FIGURE 08: Four modes of engagement . . . . . 60
- FIGURE 09: Depth of knowledge framework for writing . . . . . 64
- FIGURE 10: Human essays contribute more unique ideas than ChatGPT-generated essays . . . . . 68
- FIGURE 11: Children's data privacy not included . . . . . 89
- FIGURE 12: Data collection and usage summary of OpenAI's privacy policy as of August 15, 2025. . . . . 92
- FIGURE 13: OpenAI data policy: Data sharing and disclosure as of August 15, 2025 . . . . . 93
- FIGURE 14: The AI continuum of dependence . . . . . 106
- FIGURE 15: The AI flywheel effect on children's learning and development. . . . . 110
- FIGURE 16: Rural versus urban Internet usage . . . . . 113
- FIGURE 17: Top 12 spoken versus online languages . . . . . 115
- FIGURE 18: The instructional core that drives educational change . . . . . 121
- FIGURE 19: AI enriched versus diminished learning experiences . . . . . 122
- FIGURE 20: A framework for action. . . . . 126

# BOXES

<b>BOX 01:</b> Data sources .....	23
<b>BOX 02:</b> Common sources for generative AI .....	24
<b>BOX 03:</b> Types of AI tools and platforms study participants commonly report using .....	26
<b>BOX 04:</b> Possible solutions for expanding access to AI in underserved areas .....	37
<b>BOX 05:</b> Time savings only occur if teacher workloads remain steady .....	40
<b>BOX 06:</b> AI and early childhood development .....	55
<b>BOX 07:</b> Reducing the tendency to offload thinking to AI .....	67
<b>BOX 08:</b> Seductive design: Banal deception, sycophancy, and anthropomorphization .....	71
<b>BOX 09:</b> Emerging safeguards for children in the age of AI .....	76
<b>BOX 10:</b> Building trust around AI .....	85
<b>BOX 11:</b> Privacy and safety for student AI .....	90
<b>BOX 12:</b> Small language models: Addressing safety concerns .....	100
<b>BOX 13:</b> Addressing students' AI dependence .....	109
<b>BOX 14:</b> Data centers .....	114
<b>BOX 15:</b> Measuring the accuracy of LLMs .....	116
<b>BOX 16:</b> From policy to implementation: Addressing the AI divide .....	118
<b>BOX 17:</b> Strategies to support students' core learning capacities .....	130
<b>BOX 18:</b> Strategies to support students' knowledge and understanding, through interdisciplinary frameworks .....	132
<b>BOX 19:</b> Strategies to support students' social and civic development .....	133
<b>BOX 20:</b> Strategies to support students' motivation and agency, including through real-world learning .....	134
<b>BOX 21:</b> Teacher-tech co-design hubs .....	135

**BOX 22: Involving students and parents in AI decisionmaking and design** ..... 136

**BOX 23: Local language and community AI initiatives** ..... 137

**BOX 24: Child-friendly product design** ..... 138

**BOX 25: Diverse research approaches to study AI and education** ..... 141

**BOX 26: IFE for AI infrastructure** ..... 148

**BOX 27: Strategies for equitable distribution of AI tools** ..... 150

**BOX 28: AI safety standards for children** ..... 153

**BOX 29: Strategies to stress test AI platforms for safety** ..... 154

**BOX 30: Child-friendly procurement criteria** ..... 159

**BOX 31: Product certifications** ..... 160

**BOX 32: Family-facing AI safety information** ..... 162

**BOX 33: Student led, peer-to-peer AI discussion** ..... 163

# EXECUTIVE SUMMARY

## INTRODUCTION

In November 2022, OpenAI released its frontier Large Language Model (LLM), ChatGPT. Within five days, ChatGPT had five million users. Within two months, that number swelled to 100 million monthly active users. By August 2025, 700 million users across the globe were using ChatGPT (Mehta 2025). Many of those users are students harnessing LLMs for brainstorming, tutoring, creating, and learning; to work better and faster; and, crucially, to outsource their thinking.

Since the debut of this most common form of generative artificial intelligence, referred to here as “AI,” the education community has been debating its promises and perils. Rather than wait for a decade to conduct a postmortem on the failures and opportunities of AI, the Brookings Institution’s Center for Universal Education embarked on a yearlong global study—a premortem—on generative AI focused on answering the following two research questions:

- What are the potential negative risks that generative AI poses to the education of children and youth?
- Assuming these potential risks, what can we begin to do now to prevent them while maximizing the potential benefits of AI?

This premortem encompassed interviews, focus groups, and consultations with 505 students, teachers, parents, education leaders, and technologists across 50 countries. Our findings are based on their experiences with AI, a close review of hundreds of studies, and a Delphi panel. Given the emergent nature of the field and the need to shed light on how AI is being implemented in education, children’s experiences and those of the adults around them are the focal point. We share our insights not as the final word on AI and education, but as a first vision of where we are already heading.

## GOALS OF THIS REPORT

This report aims to help readers from many fields understand the current landscape of benefits and risks of AI in children’s education and help them identify concrete actions they can take to leverage AI for transformational educational benefits. Ultimately, the trajectory of AI will be determined not by fatalism or passive acceptance, but by the deliberate choices and sustained efforts of all of us working together to bend the arc of AI implementation toward educational experiences that help all children flourish academically, socially, and civically.

## FOUNDATIONAL PREMISES OF THIS REPORT

This report is grounded in three critical premises that shape our findings and recommendations.

First, a core premise of this study is that *children’s learning* is fueled by social relationships and their holistic development (National Scientific Council on the Developing Child 2004). Decades of research in developmental science demonstrates that children’s ability to learn—including in formal schooling—is shaped by the interconnected growth of their cognitive, social, and emotional capacities. These domains

are mutually reinforcing: language and executive-function skills support social interaction; emotional regulation influences attention and persistence; social relationships fuel motivation and meaning-making. Learning, therefore, cannot be separated from the broader developmental systems in which children grow (Bronfenbrenner 1994; Immordino-Yang and Damasio 2007; National Academies of Sciences, Engineering, and Medicine et al. 2018). Hence, we have examined here the role of AI on children’s cognitive, social, and emotional development both in and out of school, which includes their growing interaction with AI friends.

Second, *schools* serve multiple, interrelated purposes in children’s lives and in society. While they are widely understood as places for academic learning, they are also central spaces for social development, personal growth, and learning how to live in a diverse community. In school, children learn to understand others’ perspectives, collaborate, negotiate differences, and build the social and emotional skills needed to participate in diverse communities. At the same time, schools are environments where children can develop capacities that prepare them for future work, meaningful life pathways, and participation in civic life. Schools can provide stability, routine, and care; support children’s well-being; and enable parents’ labor-force participation (Winthrop and McGivney 2015). Thus, we broadly examine AI in terms of the role schools play in academic, developmental, social, civic, and economic roles — all of which are deeply intertwined in children’s everyday educational experiences.

Finally, *technology* has long been heralded as inherently “transforming” education—a means to overcome existing constraints in the education system. The prevailing narrative suggests that the very adoption and use of technology represent innovation and progress, with the implication that schools and students who do not use these tools will suffer academically, vocationally, and personally. This same rhetoric now drives the adoption and use of AI in schools (Reich 2025).

However, decades of investment and implementation have demonstrated that technology’s educational benefits have been mixed at best (UNESCO 2023). Multiple rigorous cross-national studies have shown that education systems investing heavily in technology do not necessarily experience improved teaching and learning outcomes (OECD 2015; West 2023). The mobile broadband example illustrates this pattern—while Internet expansion correlates with economic development, a study of 2.5 million 15-year-olds from 82 countries suggests that the rollout of 3G coverage from 2000-2018 produced statistically significant declines in math, reading, and science scores, as well as students’ social relationships and sense of belonging (Jain and Stemper 2024). Rhetoric arguing that technology adoption in and of itself represents innovation and progress is not only false but undermines society’s ability to discern how to effectively harness AI to advance children’s education (Reich 2025).

Research indicates that technology contributes most effectively to educational improvement when embedded within carefully designed and implemented strategies (Hardman et al. 2019). This requires several conditions. Tools and platforms must be designed ethically and responsibly, grounded in the learning sciences. Schools and families must work in partnership to ensure that children’s AI use supports—not harms—their learning and development. AI tools must support human relationships, including the teacher-student relationship, using sound pedagogical practices designed to augment, rather than substitute student learning. Finally, educators and students must remain aware of both the benefits and harms these technologies present. When these conditions align, AI has the potential to meaningfully enhance educational outcomes while minimizing risks to learners. This conceptual grounding provides valuable guidance for understanding and discussing AI in educational contexts.

## THE CHALLENGE OF ANALYSIS

Analyzing a nascent and dynamic innovation like AI presents distinct challenges. The technology's rapid evolution, its sui generis nature, and the absence of extensive precedents make establishing stable frames of reference or accurate predictions of benefits or harms difficult. We are operating within a context of scarcity of rigorous and longitudinal research evidence on the effects of AI in education on student learning and well-being. None of us, not even AI's creators, can predict its potential dangers or benefits with complete accuracy. While new developments will inevitably render some information in this report moot by the time of publication, the data presented here remains acutely relevant for this critical juncture in AI's development. Accepting uncertainty and incorporating it into our reasoning allows us to speak wisely about what we observe and to better prepare educational environments for both the opportunities and challenges that generative AI brings.

## CURRENTLY, THE RISKS OF AI OVERSHADOW THE BENEFITS

Given the trajectory of AI's implementation and use at this point in time, this report identifies two potential outcomes:

**Potential outcome 1: AI-enriched learning experiences.** In communities with access to AI, well-designed AI tools and platforms can offer students a number of learning benefits *if* deployed as a part of an overall, pedagogically sound approach. AI can also expand access to education for students who were previously excluded from teaching and learning experiences. In schooling, AI enriches learning when it expands and deepens the capabilities of—and interactions between—students, teachers, and content, which are at the heart of education. AI tools explicitly designed to support the learning of children and youth, including experiences bounded by accurate content and safety guardrails, can foster enriched learning experiences.

**Potential outcome 2: AI-diminished learning experiences.** In communities with regular access to AI, overreliance on these tools and platforms can put children and youth's fundamental learning capacity at risk. These risks can impact students' capacity to learn, their social and emotional well-being, their trusting relationships with teachers and peers, and their safety and privacy. Indiscriminate AI implementation also risks exacerbating social divides. In schooling, AI diminishes learning when it restricts and weakens the capabilities of—and interactions between—actors in the instructional core, namely students, teachers, and content. AI tools designed for the general public, with limited safeguards and a risk of inaccurate information, can lead to diminished learning experiences.

**Ultimately, we find that at this point in its trajectory, the risks of utilizing AI in education overshadow its benefits.** This is largely because the risks of AI differ in nature from its benefits—that is, these risks undermine children's foundational development. For example, when trust between students and teachers is diminished—a current risk of AI implementation—the benefits of AI-enriched teaching and learning materials cannot be fully realized. Understanding the distinction between enriched and diminished learning experiences is an important first step toward mitigating AI's risks and harnessing its benefits in education.

## **IT'S NOT TOO LATE: RECOMMENDATIONS FOR ACTION**

Above all, this report is a call to action. While AI's potential negative risks and the damages it has already caused are daunting, they are fixable. We should neither capitulate to these harms nor focus solely on limiting their repercussions.

---

**We urge all relevant actors to identify at least one recommendation to advance over the next three years.**

---

Instead, as the Section VI exhorts, governments, education funders, technology companies, education systems, civil society organizations, educators, families, caregivers, community actors, researchers and academics, and students have the agency, the capacity, and the imperative to enable students to harness AI's academic, vocational, and personal benefits while simultaneously shielding them from its potential risks and actual harms.

Thus, the final section of this report presents 12 recommendations for multiple stakeholders organized around three foundational pillars that together form a comprehensive framework for action. We urge all relevant actors to identify at least one recommendation to advance over the next three years.



**PROSPER:** Recommendations under the Prosper pillar focus on transforming teaching and learning experiences so that children and youth can thrive in an education system where AI is omnipresent. Students can *prosper* through carefully titrated AI use (knowing when to teach with and without AI, using AI only when it enhances rather than replaces student effort and cognitive engagement), high-quality pedagogical integration (combining AI with evidence-based practices that prioritize deeper learning), and collaborative design and research (co-designing tools with educators and communities while conducting rigorous research on when and how AI supports learning). These efforts will help children and youth develop the capabilities they need to flourish in an AI-driven world.

**PREPARE:** Recommendations under the Prepare pillar focus on building the knowledge, capacity, and structures needed for students, educators, families, and education systems to integrate AI ethically, effectively, and humanely. The education community can *prepare* for a rapidly evolving AI-infused world through holistic AI literacy (developing understanding about AI's capabilities, limitations, and implications), robust professional development (equipping educators with knowledge and skills to teach with and about AI), and systemic planning and access (establishing clear visions for ethical AI use while expanding equitable access). These efforts will ensure that students, educators, families, and education systems can navigate the world of AI with confidence and responsibility.

**PROTECT:** Recommendations under the Protect pillar include developing and implementing safeguards on AI for student privacy, safety, emotional well-being, and cognitive and social development. Students can be *protected* through ethical and trustworthy AI design (protections embedded into the technology during the design phase), responsible governance (strong regulatory frameworks), and adult guidance (modeling healthy technology use at home and in schools). These efforts will help ensure that AI enhances learning while keeping children safe, supported, and able to thrive in an AI-infused world.

## AI can...

DEGRADE trust in education

Put student safety at risk

Automate assessment

Reduce reading and writing skills

IMPROVE learning

Undermine family bonds

Power Intelligent Tutoring Systems

Widen linguistic and cultural gaps

Enable cognitive offloading

IMPEDE social and emotional development

Support student reading

Increase surveillance of students

OPTIMIZE teachers' time

Expand socioeconomic and regional inequity

Support underserved and neurodivergent students

Encourage harmful behavior

DEEPEN equity divides

**AI can...**

EXPAND access

Displace trust in humans

Weaken critical thinking and creativity

Enhance assistive technologies

ERODE students' autonomy and agency

Damage teacher-student relationships

Support individualized attention

ADVANCE assessment

Exacerbate mental health concerns

Manipulate emotions

Undermine democracy and freedom

THREATEN safety

Help students with writing

Share private student data

Increase loneliness

Expose children to sexualized content and harassment

EXTEND learning

UNDERMINE cognitive development

Identify at-risk students



## I INTRODUCTION

*A new technology sometimes creates more than it destroys. Sometimes it destroys more than it creates. But it is never one-sided.*  
— Neil Postman (1990)

What role will artificial intelligence (AI) play in shaping the future of education? Will it transform education, or will it generate harms that undercut these gains? As educators contemplate integrating AI into teaching and learning, how can they embrace its potential while preserving human agency, deep learning, and pedagogical judgment? How can we navigate and manage the unexpected and unintended consequences of the relationship between human cognition, institutional imperatives, and AI?

As the above questions imply, generative artificial intelligence (hereafter referred to in this report as “AI”) presents opportunities and threats, promising unprecedented innovation—and disruption. This duality prompted the Center for Universal Education at the Brookings Institution to examine AI in relation to the education of children and youth. This is a distinctly urgent global phenomenon: children—whom we define as those under 18—access AI products regardless of national borders. And unlike the use of AI by adult professionals, when children use AI, it affects their cognitive and emotional development.

After OpenAI released ChatGPT, its frontier Large Language Model (LLM), in November 2022, AI began to shape how children learn on an international scale. In crisis-affected communities lacking traditional schools, children engage with adaptive lessons through messaging platforms or AI tutors. Children with speech and communication challenges use AI to create digital versions of their own voices in order to communicate with classmates and teachers. Around the world, children interact with robotic teacher assistants and use AI-embedded Virtual Reality (VR) to immerse students in real-world learning environments (Chowdhury 2024; ANI News 2022).

Students' AI use extends beyond school, blurring the lines between entertainment, social connection, and academic work. They talk to AI "friends" via social media, confide in chatbots for emotional support and relationship advice, and outsource schoolwork to AI (and then employ AI humanizers to disguise their AI-generated homework). For some students, AI use is a lifeline to a larger world and a fuller sense of self (Franze et al. 2023). For others, it represents a dangerous slide toward dependence, school withdrawal, social isolation, and, tragically at times, self-harm (Common Sense Media 2025a; Duffy 2024; Duffy 2025; McLennan 2025).

These examples suggest that concerns about AI's duality and disruptive nature are not unfounded. The impacts of new technologies are often ecological in scope, affecting not one thing but *everything*, including human behavior (Postman 1995). Every technological advantage yields corresponding, and often unforeseen, disadvantages whose effects are spread unevenly (Postman 1995; Kranzberg 1986). Social media exemplifies this duality. Its verbal immediacy and open, networked structure have fostered unprecedented opportunities for human creativity, communication, and collaboration (Burns and Bodrogini 2011). Yet these same platforms serve as vectors for social maladies, scaling misinformation, amplifying polarization, and influencing users' mental health (OECD 2021a; Twenge et al. 2021;

Haidt 2024). Financial incentives to maximize user engagement often exacerbate these harms.

As this report will document, the adoption and integration of AI tools presents a complex landscape of educational benefits and risks, disruption, and misalignment that are potentially more consequential than those experienced with social media. When employed inventively, new technologies can be transformative. Adaptive AI learning tools can personalize instruction and enhance student agency; AI can reduce administrative burdens on teachers, freeing time for meaningful interaction; automated tutoring can increase student motivation and support differentiated learning. However, AI tools can also prompt overreliance, emotional and cognitive dependence, and diminished critical thinking. This duality demands that policymakers, decisionmakers, and educators approach educational technology implementation with a sophisticated understanding of both its transformative potential and its inherent risks to student well-being and learning.

AI stands as a remarkable testament to human imagination and innovation; its tools and systems embody human values and priorities, promising enormous benefits to human beings (PBS: Firing Line with Margaret Hoover 2025; Quezzaire 2025). Its stakes for education are particularly high because AI's influence will increasingly extend to society's most vulnerable population—our children and youth.

## **This report: Its purpose and main message**

This report is designed to serve as a resource for educational leaders, government officials, technologists, school administrators, teacher and family organizations, student networks, funders, and other stakeholders responsible for developing

---

**We worry that the larger education community may miss the forest for the trees: A narrow focus on developing effective AI-supported teaching approaches could obscure how students' very ability to learn is being undermined by AI overuse, inappropriate use, and non-productive use, both in and, increasingly, outside the classroom.**

---

and implementing AI in education. The focus is squarely on AI and children and youth's educational experiences, specifically on students attending primary and secondary school.

The information presented here represents a snapshot of AI experiences during the early stages of its adoption, integration, and use in education, both inside and outside the classroom. These perspectives illuminate both the issues at stake and the responses they provoke among practitioners navigating this emerging landscape. Given the rapidly changing nature of this field, they are by no means definitive.

Our research identifies potential outcomes if AI deployment continues on its current course. Benefits include enriched learning experiences through expanded access to education; improving teaching and learning through personalized instruction; providing multimodal ways of accessing lessons; and allowing educators more time with students.

Yet, as this report will show, these benefits currently pale in comparison to the risks of AI. For our study participants, these risks include undermining children's cognitive, social, and emotional development; reducing trust between students, teachers, and families; threatening study privacy; and opening a new AI equity divide. We worry that the larger education community may miss the forest for the trees: a narrow focus on developing effective AI-supported teaching approaches could obscure how students' very

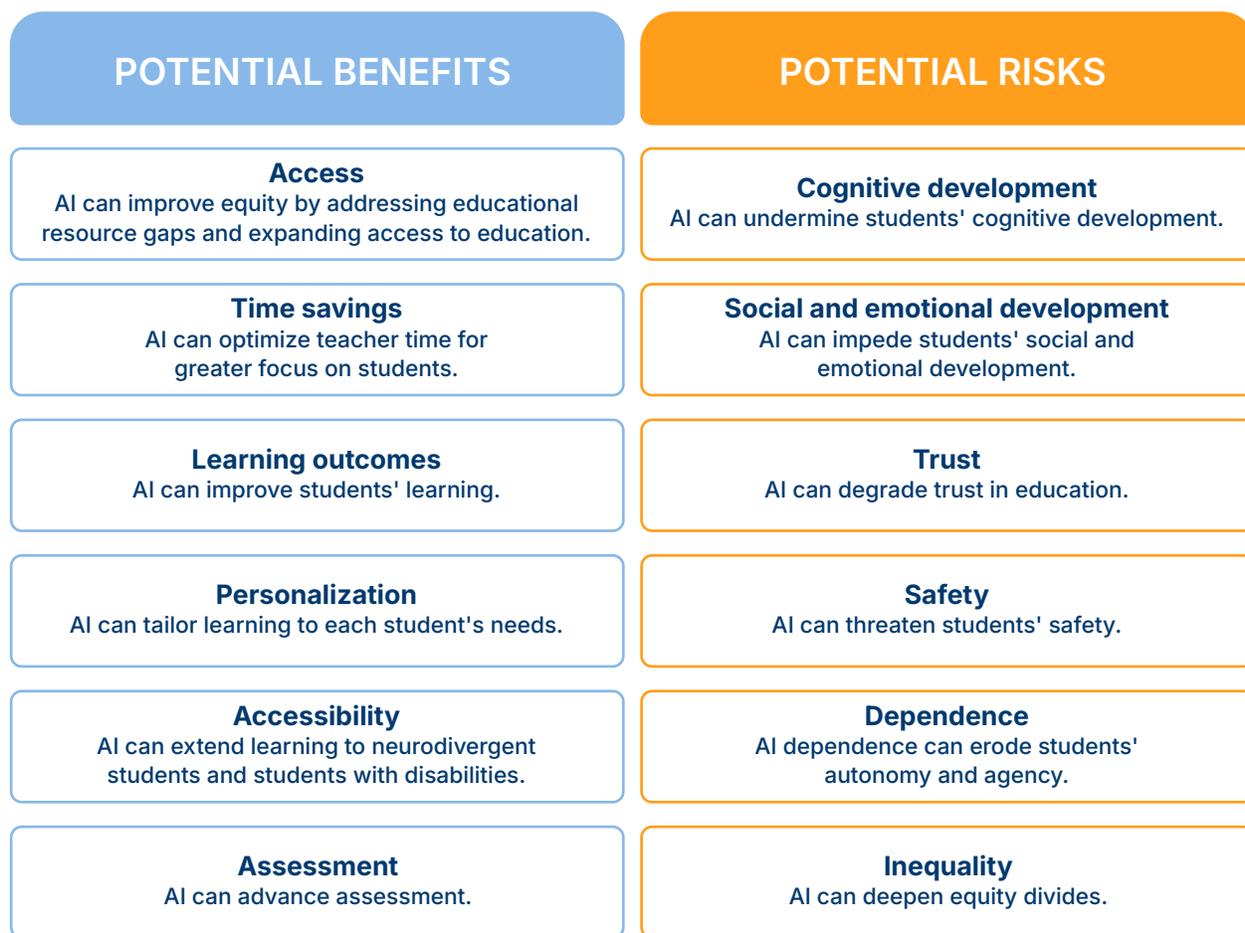
ability to learn is being undermined by AI overuse, inappropriate use, and non-productive use, both in and, increasingly, outside the classroom. The qualitatively different nature of these risks—which undermine the cognitive and social development essential for successful teaching and learning—demands that we establish frameworks to harness AI's transformative potential with safeguards.

All technologies create both hoped-for benefits and unanticipated negative consequences, often disrupting systems in ways their developers never imagined. Recognizing this dynamic allows us to invert traditional discussions about AI from such passive queries as, "What will AI do to us or our future?" to active, agency-driven questions like, "What will we do to create the educational future we want for our children" and "How will we control AI?" This reframing helps us understand that AI's negative outcomes in education are not predetermined. As human beings with autonomy, we can approach AI use in education with intentionality and discernment so that our technological future does not resemble our technological past.

To successfully bend the arc of AI implementation toward improved student learning, we must mitigate risks and harness benefits according to three "pillars": prosper, protect, and prepare. All three pillars must work in concert to shift AI's course so that all children and youth may access an education that helps them flourish as humans in a technological world. We can utilize AI's power to help children prosper educationally by using AI

FIGURE 1

## Potential benefits and risks of AI in education



© Center for Universal Education at the Brookings Institution, 2026.

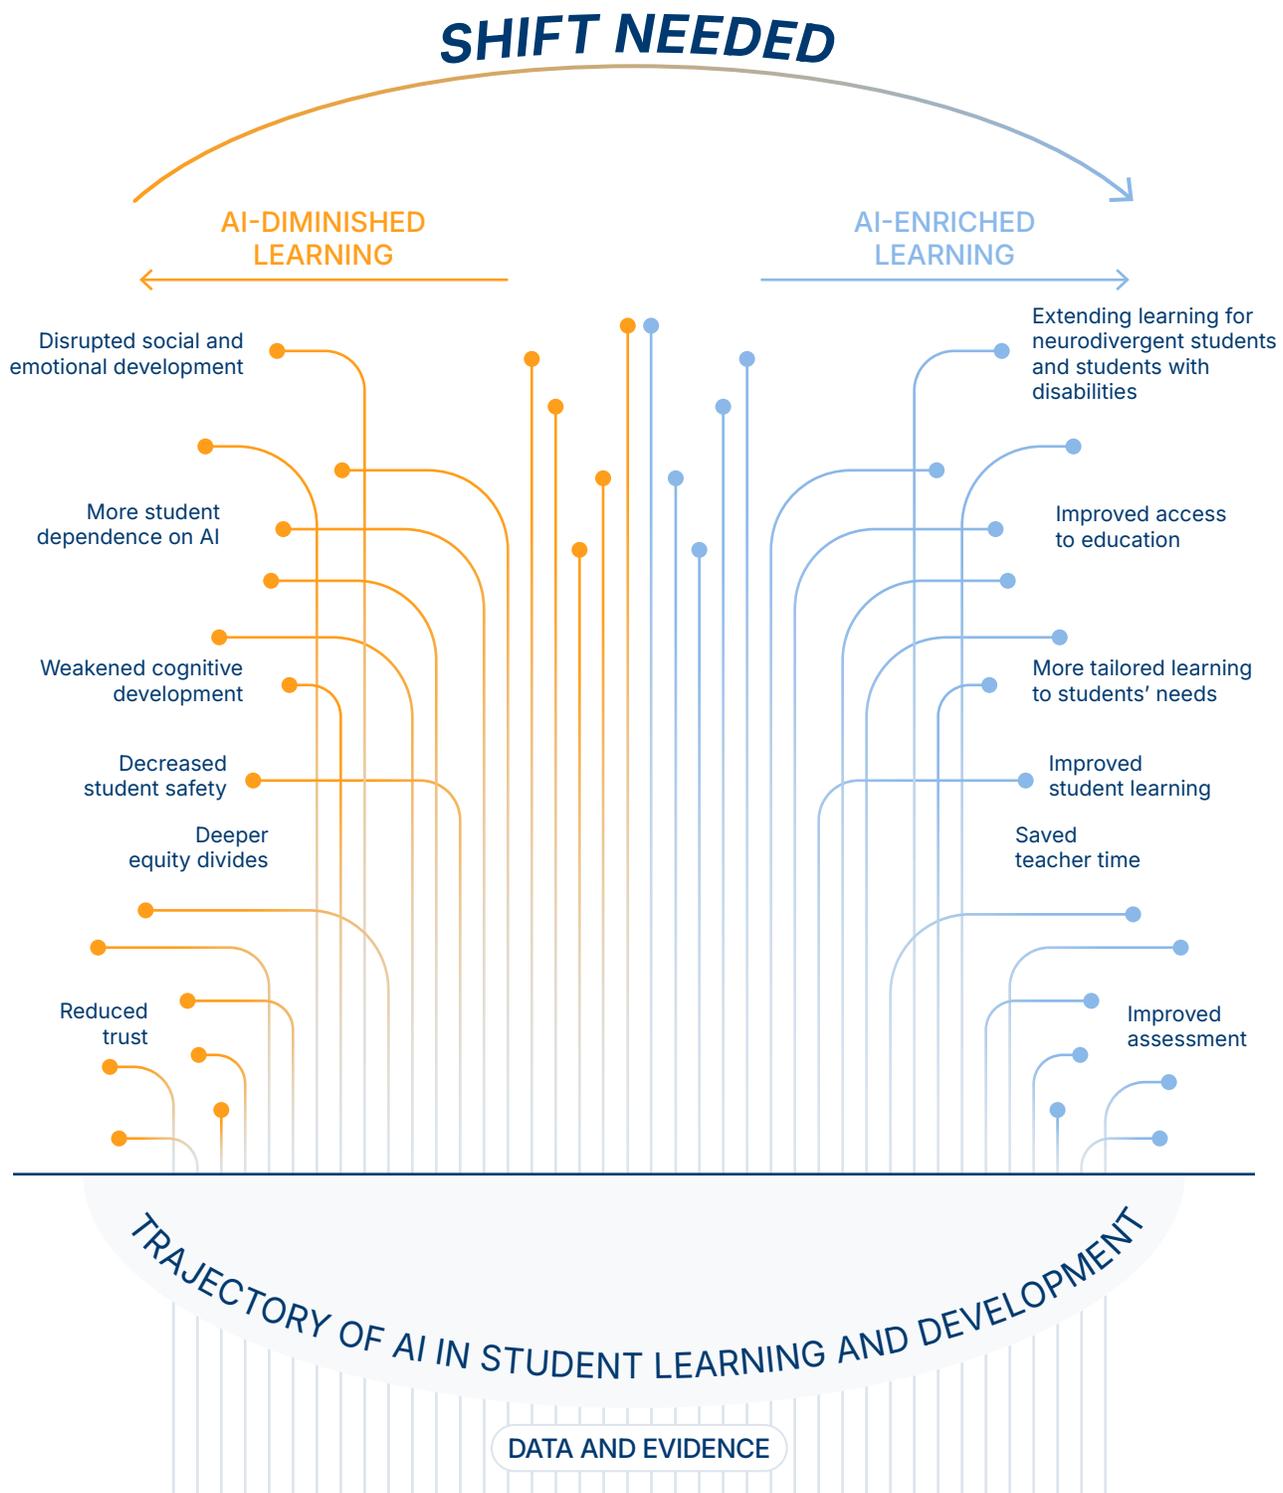
tools that embed research-proven instructional design and learning sciences as part of pedagogies that use AI to complement and strengthen, not replace, human knowledge, skills, and judgment, while avoiding AI deployment when it undermines essential developmental processes such as self-regulation, metacognitive development, or critical student-teacher relationships (Cukurova 2024).

We can protect students' privacy, safety, emotional well-being, and cognitive and social development by ensuring safe access to ethical and responsible AI tools that support strong in-person relationships and productive struggle. We can help educators, families, and students prepare for AI by building the knowledge, capacity, and structures needed

to integrate AI effectively, while reflecting on the ethical implications and creative uses of AI and the ways technology shapes our behavior, our social interactions, and our relationship with knowledge. Ultimately, our hope is that our findings and recommendations help shift the course of AI's progression so that all children and youth everywhere may access learning that helps them develop, grow, and thrive as human beings in an AI-infused world.

FIGURE 2

Bending the arc from diminished to enriched AI learning experiences



© Center for Universal Education at the Brookings Institution, 2026.

# Foundational premises of this report

This report is grounded in three critical premises that shape our findings and recommendations.

First, a core premise of this study is that *children's learning* is fueled by social relationships and their holistic development (National Scientific Council on the Developing Child 2004). Decades of research in developmental science demonstrates that children's ability to learn—including in formal schooling—is shaped by the interconnected growth of their cognitive, social, and emotional capacities. These domains are mutually reinforcing: language and executive-function skills support social interaction; emotional regulation influences attention and persistence; social relationships fuel motivation and meaning-making. Learning, therefore, cannot be separated from the broader developmental systems in which children grow (Bronfenbrenner 1994; Immordino-Yang and Damasio 2007; National Academies of Sciences, Engineering, and Medicine et al. 2018). Hence, we have examined here the role of AI on children's cognitive, social, and emotional development both in and out of school, which includes their growing interaction with AI friends.

Second, *schools* serve multiple, interrelated purposes in children's lives and in society. While they are widely understood as places for academic learning, they are also central spaces for social development, personal growth, and learning how to live in a diverse community. In school, children learn to understand others' perspectives, collaborate, negotiate differences, and build the social and emotional skills needed to participate in diverse communities. At the same time, schools are environments where children can develop capacities that prepare them for future work, meaningful life pathways, and participation in civic life. Schools can provide stability, routine, and care; support children's well-being; and enable parents'

labor-force participation (Winthrop et al. 2015). Thus, we broadly examine AI in terms of the role schools play in academic, developmental, social, civic, and economic roles — all of which are deeply intertwined in children's everyday educational experiences.

Finally, *technology* has long been heralded as inherently "transforming" education—a means to overcome existing constraints in the education system. The prevailing narrative suggests that the very adoption and use of technology represent innovation and progress, with the implication that schools and students who do not use these tools will suffer academically, vocationally, and personally. This same rhetoric now drives the adoption and use of AI in schools (Reich 2025).

However, decades of investment and implementation have demonstrated that technology's educational benefits have been mixed at best (UNESCO 2023). Multiple rigorous cross-national studies have shown that education systems investing heavily in technology do not necessarily experience improved teaching and learning outcomes (OECD 2015; West 2023). The mobile broadband example illustrates this pattern—while Internet expansion correlates with economic development, a study of 2.5 million 15-year-olds from 82 countries suggests that the rollout of 3G coverage from 2000-2018 produced statistically significant declines in math, reading, and science scores, as well as students' social relationships and sense of belonging (Jain and Stemper 2024). Rhetoric arguing that technology adoption in and of itself represents innovation and progress is not only false but undermines society's ability to discern how to effectively harness AI to advance children's education (Reich 2025).

Research indicates that technology contributes most effectively to educational improvement when embedded within carefully designed and implemented strategies (Hardman et al. 2019). This requires several conditions. Tools and platforms must be designed ethically and responsibly, grounded in the learning sciences. Schools and

families must work in partnership to ensure that children’s AI use supports—not harms—their learning and development. AI tools must support human relationships, including the teacher-student relationship, using sound pedagogical practices designed to augment, rather than substitute student learning. Finally, educators and students must remain aware of both the benefits and harms these technologies present. When these conditions align, AI has the potential to meaningfully enhance educational outcomes while minimizing risks to learners. This conceptual grounding provides valuable guidance for understanding and discussing AI in educational contexts.

## Our methods: A premortem on AI and children’s education

In September 2024, the Center for Universal Education launched the Brookings Global Task Force on AI and Education with the goal of conducting a premortem on AI and education. A premortem anticipates potential failures before implementing an innovation. Unlike a postmortem, which examines failure after occurrence, the “prospective insight” gained from a premortem helps to anticipate potential issues with any innovation so they can be corrected or mitigated before norms of use become entrenched (Burns and Winthrop 2024; Mitchell et al. 1989; Veinott et al. 2010). Rather than waiting years to discover the negative impacts of AI, our goal is to help educators, parents, policymakers, technology companies, and students themselves identify steps they can take now to prevent potential negative impacts and support positive outcomes. Importantly, a premortem is not an exercise in pessimism; we expect AI can enrich students’ learning experiences. However, to harness these benefits, harms and unintended consequences

require deliberate anticipation to prevent or mitigate.

The two guiding research questions of this report are:

- What are the potential negative risks that generative AI poses to the education of children and youth?
- Assuming these potential risks, what can we begin to do now to prevent them while maximizing the potential benefits of AI?

This report focuses primarily on school-aged children and youth (defined in this report as those under 18) and the ways AI is shaping their education. Our student study participants ranged from ages 13 to 17. We begin at age 13 since this is the threshold at which AI companies, and other technology companies, commonly permit legal access to AI products. However, as will become evident throughout this report, many students younger than 13 are already accessing these products.

From September 2024 to October 2025, we examined AI’s evolving role in education through focus groups, interviews, and expert consultations with over 500 educators, policymakers, technologists, civil society leaders, task force steering committee members, educators, students, and parents across 50 countries, alongside a literature review of over 400 articles and guidance from a Delphi panel of 21 experts (See Box 1). Our understanding has been enriched by these diverse perspectives and throughout the report we refer to those participating in interviews, focus groups, and consultations as *study participants*. At the heart of our inquiry are the voices of those on the frontlines: teachers, students, and parents and caregivers. Their ideas and experiences form the foundation of this report.

A more detailed explanation of our research, including the data frame, data collection instruments, and analysis procedures, is available in Annex C and on the Brookings website. While

a premortem approach cannot provide definitive answers to AI's potential risks and benefits—and given the rapidly evolving nature of AI, many answers remain unknowable—it is essential that we interrogate new technologies even in the face

of such uncertainty (Harari 2024). A premortem approach helps us better understand the range of issues we must attend to in order to support all children's learning and development.

## BOX 1

# Data sources

**Interviews and focus groups:** To capture the emerging nature of AI in education, we conducted group interviews and focus groups. The perspectives and lived experiences of these key informants form the foundation of our findings. **Group interview** participants included students, educators, and parents and caregivers, hereafter referred to as parents. **Focus groups** included a broad range of key informants we call in this report “experts”: university researchers, education leaders, AI education practitioners, government officials, members of think tanks and non-governmental organizations, and business and community leaders.

**Review of the research on AI in education:** Guided by the two research questions, and to better situate study participants' experiences within a theoretical context, we conducted an extensive review of studies on AI in education published after November 2022. Given the nascent state of this field, the rapid pace of AI development, and the scarcity of experimental and longitudinal research, the review adopted a pragmatic approach, also incorporating preliminary findings and smaller-scale investigations, policy documents, essays, and case studies. We reviewed over 400 articles, using these insights to inform our analysis and findings.

**Delphi panel:** The Delphi approach uses a series of open-ended questions posed to highly knowledgeable experts within a structured consensus-building approach. This makes it particularly useful for predicting trends and forming guidance when clear data or field consensus does not yet exist (Khodyakov 2023). We convened a Delphi panel of education leaders and AI specialists who analyzed and commented on the combined data collected from study participants and the literature review. However, unlike the key informants mentioned above, Delphi panelists did not participate in interviews. The views and guidance of Delphi panelists are integrated throughout this report.

**Consultations:** We conducted consultations with a wide variety of stakeholders in order to inform each stage of the research process—getting feedback on the methodology, stress testing data collection instruments, surfacing gaps in preliminary findings, and final data analysis. The participants in these consultations ranged from high-level decisionmakers (including those on the task force steering committee) to funders and practitioners such as ed-tech leaders, nongovernmental organization representatives, child development specialists, learning scientists, and technologists.

FIGURE 3

## Map of study participants' countries



© Center for Universal Education at the Brookings Institution, 2026.

## What is generative AI?

AI has a long history, including in education. Predictive AI, for example, has long been used in intelligent tutoring systems (ITSs) to guide students in their formulation of a correct response to a test question.

Generative AI, the focus of this report, refers to a class of technologies that generate novel outputs such as text, images, audio, code, or video in response to user prompts. These technologies include LLMs, diffusion models, and other neural network architectures designed for content generation. Generative AI is a form of machine learning that creates new content rather than simply analyzing or classifying existing data. These systems are typically trained on vast datasets composed of publicly available digital content—words, images, and sounds (see Box 2 for

common sources for generative AI). Contemporary models increasingly incorporate curated, domain-specific, and proprietary data tailored to specific applications, including educational ones.

### BOX 2

## Common sources for generative AI

- Common Crawl (a vast publicly available dataset across more than 40 languages)
- WebText2 (webpages linked from Reddit posts with at least three upvotes)
- Books1 and Books2 (collections of online books, such as Project Gutenberg's public domain texts and self-published works from Book Corpus)
- English-language Wikipedia (Rettberg 2022)

AI is not a single tool. Rather, within educational contexts, it encompasses a diverse and often overlapping ecosystem of platforms and applications designed for different purposes. They range from *education-specific* AI products used directly by students and created by commercial technology companies to the incorporation of AI into *teacher-facing products* created by long-standing education technology companies. AI also includes *learning products* on video game platforms and the AI “*friends*” or “*companions*” that students access through social media platforms. Education technology providers are typically for-profit companies, reflecting the substantial growth in the ed-tech sector (Grand View Research 2025), although such providers may also include universities or governments. In South Korea, for example, the Korea Education and Research Information Service, a government educational organization leading technology development in Korean schools, created the AI courseware platform used across the nation’s schools (KERIS 2025).

Box 3 broadly categorizes the seven types of AI tools that study participants reported using. Throughout this report, we will, where data permits, make explicit distinctions among the various forms of AI under discussion. For example, we will endeavor to distinguish between students engaging in open-ended unmediated discussion with LLMs versus students using creative AI technologies like video or image generators or teachers using AI-supported adaptive learning programs with vetted content. Yet the vast majority of study participants report using LLMs (specifically, ChatGPT) for education activities; as reported, teacher use of AI tools outpaces that of students.

Although AI systems generate responses that often appear intelligent and contextually appropriate, their outputs result from sophisticated statistical pattern-matching rather than genuine knowledge or creative reasoning (Kleiman and Gallagher 2024). AI systems and human intelligence are completely different entities, a distinction that carries profound implications for educational practice. While AI is far ahead of humans in computations, at the time of

writing it currently remains less proficient in skills like cognitive flexibility and long-term planning in novel situations (Cukurova 2024, 471). AI represents knowledge stripped of experience and cannot replace the wide range of human experiences and interactions (Kleiman and Gallagher 2024). As we consider bringing AI into learning environments, it is important for all those involved in education to keep the above facts in mind to ensure that we maintain clarity about which uniquely human capacities we seek to develop in our students.

Understanding AI is one part of the larger focus of this report; equally important is how and where children and youth use these tools. We turn now to this discussion.

### BOX 3

## Types of AI tools and platforms study participants commonly report using

- **Teacher productivity and planning tools:** Platforms such as *Gradescope*, *MagicSchool*, and *Notion.ai* support teachers' administrative and instructional workflows. They automate grading and feedback, generate standards-aligned lesson and assessment materials, produce differentiated resources, and assist with documentation such as individualized education plans and progress reports.
- **Student-facing adaptive and learning systems:** AI-powered platforms like *MATHia*, *Knewton Alta*, and *DreamBox* personalize learning by analyzing student responses in real time, adjusting content difficulty, identifying misconceptions, and providing targeted feedback to close knowledge gaps.
- **General-purpose generative AI:** LLMs such as *ChatGPT*, *Claude*, and *Gemini* function as versatile tools for both teachers and students. These tools are not designed to teach; they are designed to be helpful task replacement tools prioritizing efficiency (Quezzaire 2025; Kestin et al. 2025). Teachers use them to draft lesson plans, design assessments, and create instructional materials. Students employ them for research assistance, writing support, coding help, and conceptual understanding across disciplines. LLMs vary in the design features used to encourage extended use and engagement on their platforms.
- **Companion platforms:** Models such as *Character.ai* and *Replika.ai* provide users with AI "companions" or "friends" that simulate sustained conversation or interaction with users. They are designed to maximize engagement and extend the time users spend on the platform.
- **Specialized academic and domain tools:** Subject-specific platforms leverage AI to support mastery in particular domains—for example, *Duolingo* for language acquisition, *Photomath* for mathematical problem solving, and *GitHub Copilot* for programming and computational tasks.
- **Institutional and administrative AI systems:** These tools and platforms are used by schools or ministries for predictive analytics, attendance tracking, or early warning systems, such as *Civitas Learning*, *PowerSchool*, or *PraxiSchool*.
- **AI-enhanced assistive technologies:** These tools support accessibility, such as speech-to-text (*Otter.ai*), text-to-speech, or real-time translation, such as *Microsoft Translator* and *Seeing AI*.



## II **CONTEXT:**

Students are increasingly accessing AI outside of school, and families continue to be on the front lines

*Our students are not regularly accessing AI but our kids are.*  
— Steering Committee member

*Parents need to be educated too. We don't know much about this technology, and if our kids are using it, we should know what it is. If we know that [AI] is beneficial for children, then by all means they should know more about it. But if it causes a loss, we should know about that, too.*  
— Parent

*A lot of schools, including mine, blocked ChatGPT, but people will go on their phones and use Snapchat AI or Meta AI. People use Gemini, they use DeepSeek. Really, there are so many AIs that even if a school bans them, people are still going to be able to access it.*  
— Student

*AI is here to stay. Kids need to explore the positives on how to use it.*  
— Parent

Young people in many parts of the globe are accessing AI through multiple pathways. For those with devices and connectivity, AI may be almost impossible to avoid. While pedagogically sound AI tools that support children's learning through vetted content are only one of the ways that students encounter AI, they are far from the only way. Many students engage with AI through school, but as many (if not more) may be regularly accessing AI outside of school for an increasing range of everyday activities. AI is ever more ubiquitous, embedded in search results, social media feeds, video games, and even toys (Zhou 2025). This reality places families, parents, and community organizations on the front line alongside teachers and school leaders as essential figures helping guide students toward effective and ethical AI use.

Any assessment of the potential impacts of AI on students' learning and development must also include a holistic analysis of how children encounter AI both in and outside of school. Hence, any strategy to mitigate AI's risks and harness its benefits needs to include a focus on supporting families and community members.

## **Widespread AI use by children and youth**

As mentioned above, multiple factors, including access to connectivity and devices, create variation in AI usage by children and youth around the world. Recent surveys across North America, Europe, Asia, and Latin America show high levels of exposure to and use of AI.

Figure 4 below reviews a selection of nationally representative surveys conducted in 2024 and 2025 asking about the AI use of children and youth. While the data are not directly comparable, given the variation in the surveys in each country, they do provide a snapshot of young people's interaction with AI in middle- and high-income countries.

**FIGURE 4**

**Children’s AI use around the world**

COUNTRY	AGES	SURVEY FINDINGS	YEAR
ARGENTINA	9–17	58% of respondents have used ChatGPT	2024
AUSTRALIA	14–17	94% of respondents have used a generative AI tool (in any capacity)	2025
SINGAPORE	13–17	80% of respondents use generative AI at least once a week	2024
SOUTH KOREA	13–19	95% of respondents have used ChatGPT	2024
TAIWAN	Junior and senior high school students	58% of respondents use AI at least once a week	2025
UNITED KINGDOM	13–18	67% of respondents have used generative AI	2025
UNITED STATES	13–17	91% of respondents used generative AI in their personal life	2025

Sources: Ipsos, 2024; Madden et al., 2024; COX Mobile, 2025; Picton et al., 2025; Centre for Evidence and Implementation, 2025; Denejkina, 2025; Child Welfare League Foundation, 2025; UNICEF-UNESCO, 2025; Joon-hyun, 2025.

While children’s use of AI will likely grow, this trend tells us little about how they access and use it, a topic we will explore next.

## The blurred lines of children’s AI use

The versatility of commercial general-purpose chatbots and the seamless integration of AI across multiple digital platforms means the lines between education technology, friendship, information query, and entertainment are blurring. Students in our study report using AI through general purpose LLMs, social media, Internet search, video games, and education technology, often moving fluidly between these platforms.

Young people use the game-based learning hub in Roblox, a gaming platform with 111.8 million daily active users across 180 countries, to hone their numeracy, science, and coding skills (Nordicity 2025; Roblox 2025a). These games are popular; one Science, Technology, Engineering, and Mathematics (STEM)-focused game created by Sesame Workshop has received 70 million visits (Roblox 2025b). Some students strategically upload class notes to Google NotebookLM to create flash cards to help with test preparation while binge-watching humorous YouTube videos. Meanwhile, others integrate AI into daily routines via social media. Snapchat’s My AI chatbot was directly integrated into the existing Snapchat app for all users in April 2023, immediately distributing AI to its nearly 150 million users aged 13–17 globally (Dixon 2025; Snapchat 2023; Snapchat 2025). Students in our study talked about sending friends humorous pictures on Snapchat before turning to

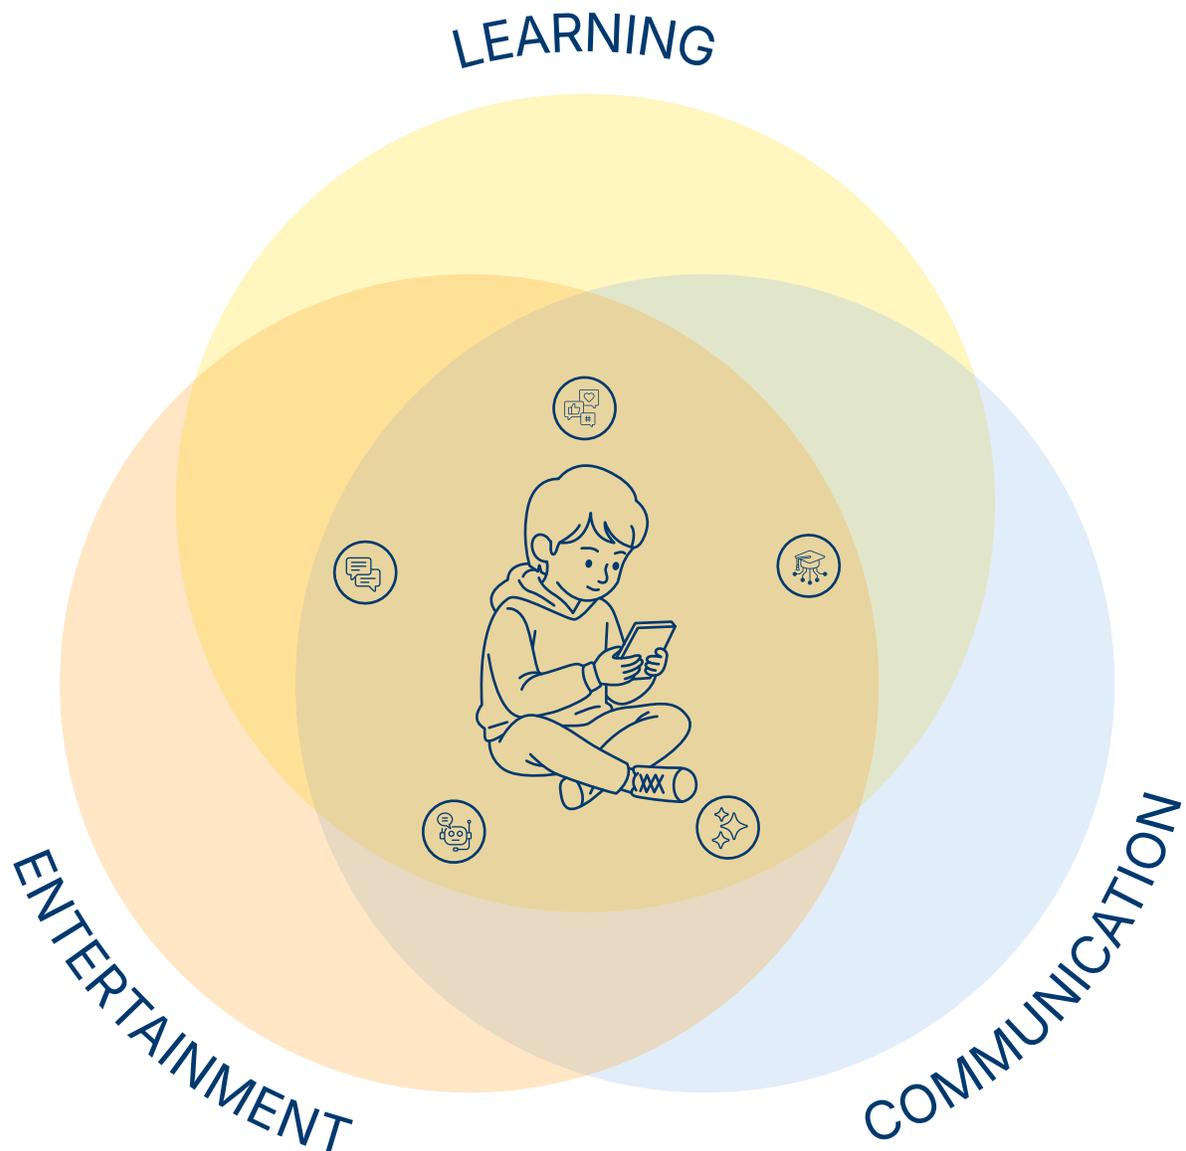
the platform’s AI feature for dating advice—and to complete their math homework for them (see Figure 5).

In our study, students reported accessing AI through a combination of platforms, including LLMs, social media platforms, Internet searches,

AI “friends,” and video games. A nationally representative survey of student AI use in the United Kingdom (UK) similarly found that young people use a mix of platforms, with LLMs like ChatGPT and Gemini and social media like Snapchat topping the list (Internet Matters 2025).

**FIGURE 5**

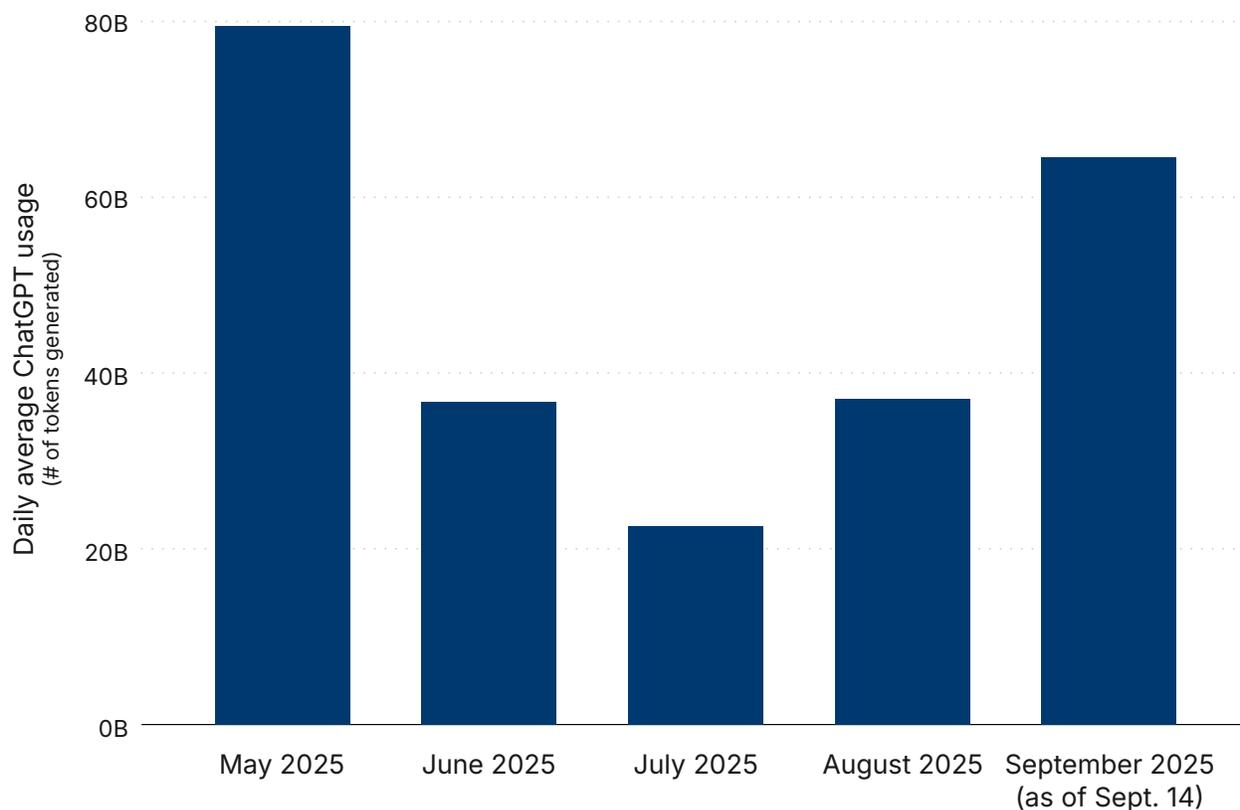
### The blurred lines of children’s AI use



© Center for Universal Education at the Brookings Institution, 2026.

FIGURE 6

## ChatGPT usage by month



Source: Wilkins, 2025; OpenRouter, 2025

It does appear that many users are using these AI tools as part of their educational experience. While AI usage data is constantly changing, worldwide data on ChatGPT usage in 2025 shows a reduction in use during June through August when many northern hemisphere schools are out of session, indicating heavy student educational use (Wilkins 2025).

## Families on the front lines

In navigating the terrain of children’s multifaceted AI use outside of school, families and community members find themselves on the front lines. In many ways, this adds an extra and unfair burden on

already overstretched caregivers. But as unfair as it is, it doesn’t take away the crucial role that these adults play in guiding children and youth to live well in an AI world.

One of the main insights from our study was the limited awareness parents and caregivers possess concerning the range of benefits and risks that AI poses to their children’s learning and development. Families are busy juggling multiple responsibilities, and often parents and caregivers haven’t had the opportunity to understand AI and its effects on children. Little support is provided in terms of programs, outreach, or AI literacy for families and community organizations who care for children outside of school, yet the majority of families we interviewed want this support.

---

## AI is helping my child by “becoming independent in their learning process...when they are stuck, they can easily just know where to get help from.”

---

Of the parents interviewed for this study, nearly all held mixed feelings about AI. On the one hand, they want their children to have the exposure and necessary skills to successfully participate in a workplace where AI will be ubiquitous. Many parents also appreciate the academic assistance AI can provide, especially when their children are stuck with schoolwork and need help. Caregivers in Ghana and Nigeria felt that AI was better able to explain schoolwork concepts and it helped children learn without waiting for adult assistance. As one parent explained, AI is helping her child by “becoming independent in their learning process... when they are stuck, they can easily just know where to get help from.”

These sentiments are pervasive in our study and across research on caregivers generally. For some parents in China, AI assists their children with learning while parents work and manage the needs of daily life. One researcher noted the rise in China of “AI babysitting, where parent circles share resources on prompting AI to guide learning” (Yan and Heting 2025). In the United States (U.S.), a 2024 national survey found that 69% of parents believed AI chatbots are a valuable tool to help their children learn more faster (Impact Research 2024).

Parents in our study also express significant concerns about AI’s potential impact on their children, concerns that are supported by broader research on parental attitudes toward AI (Internet Matters Team 2025). The main concerns center on the harms of AI vis-à-vis children’s cognitive abilities and their emotional well-being because of overuse.

As we will explore further, and discuss in the recommendations in Section VI, the complex interplay between AI benefits and risks on children’s learning and development will require that we engage families as active partners with teachers and schools.



## III POTENTIAL OUTCOME 1: AI can enrich learning if well-designed and anchored in sound pedagogy

When ChatGPT burst onto the scene in November 2022, many schools approached its adoption with an abundance of caution about its novelty, as well as its impact on student safety and learning. A number of education systems and schools initially banned its use (Johnson 2023). Yet within a matter of months, bans gave way to cautious experimentation and, in many cases, enthusiastic adoption across multiple education systems globally (Lake 2023).

AI has numerous design elements that positively impact its adoption, use, and benefits. Unlike specialized technologies that often require expensive licenses or infrastructure, many LLMs remain free, though with use restrictions and functionality constraints. Their power rests in their generality; they can produce a variety of creative outputs including musical compositions, sonnets, novels, or videos. A range of AI powered creativity tools—from vibe coding to video production—provide students and educators with new

ways to explore ideas and bring them to life. These outputs can be refined and made more specific with increasingly targeted user prompts (Burns 2024, 53). This refinement capability means that even schools without access to specialized curriculum resources, local multimedia content, or specialized staff can adapt AI tools to meet their specific instructional needs. Over time, new and varied uses of AI that enhance student learning will likely emerge as the technology advances, and educators become more adept at using it.

AI tools are almost intuitive in their use. Many dedicated AI platforms, like ChatGPT, DeepSeek, or Khanmigo, involve user prompts instead of confusing menus. AI is embedded in common technology tools, like Google and Microsoft's productivity suites, allowing users to benefit from AI-powered features without needing to learn new skills. These characteristics—rapid adoption, accessibility, generality, the ability to iteratively refine outputs, and ease of use—create the foundation for the educational benefits that follow, with particular promise for democratizing access to powerful learning tools across schools regardless of resource levels.

According to study participants, AI offers benefits for students and their teachers that fall into two distinct, though interrelated, categories. First, for teachers, AI functions as a powerful productivity tool that replaces routine tasks—lesson planning, grading, assessment design, administrative work—freeing them to focus on high-yield activities such as differentiated support, individualized feedback, and relationship building. Second, AI offers personalized learning benefits directly to students through adaptive platforms, writing tools, and tutoring programs that tailor instruction to individual needs. Thus, as will be discussed in this section, AI's dual capacity—to free teacher time through task automation and to provide personalized learning tools directly to students—can, in the words of one teacher, “democratizes access” to individualized learning and support, enabling teachers to focus their limited time on students who need the most intensive supports.

## Benefit 1: AI can improve equity by addressing educational resource gaps and expanding access to education

*AI represents the democratization of academic assistance.*

— Expert

*AI makes learning better for the students. We are not exposed to knowledge like in other countries. With AI, we can catch up just like learners in other countries.*

— Student

One of the most compelling arguments for AI in education is its potential to advance educational equity through expanding access to educational opportunities, especially for students who have previously been excluded from formal schooling. Technology has a long history of providing access to learning for communities lacking quality in-person teaching. Interactive Audio Instruction, live radio tutoring, and television classes like Mexico's *Telesecundaria* + program can scale learning opportunities; deliver national curriculum to learners in remote areas; support community volunteer teachers who have little or no training; provide targeted instruction to out-of-school youth and refugees; and offer some degree of standardized educational quality where qualified teachers are absent (Burns 2023b, 38–39).

The need for such interventions is urgent. Globally, millions of children struggle just to access education, including 251 million out-of-school children (O'Donnell 2025, UNESCO 2024). Students

in resource-limited communities worldwide have the most to gain from high-quality education yet are often taught by poorly qualified teachers (UNESCO Institute for Statistics and Education for All Global Monitoring Report 2014; Qin and Bowen 2019; UNESCO and International Task Force on Teachers for Education 2024). Technology, including AI, can help bridge this capacity gap by providing access to quality instructional materials where human expertise is absent or limited (Wang et al. 2025).

---

### **In Afghanistan, where the Taliban has banned women from post-primary education, girls and women have harnessed the potential of the Internet and AI to further their education.**

---

A lack of access to education doesn't simply mean exclusion from learning; it often represents isolation from a social and educational community. For those outside formal education systems—religious and ethnic minorities, internally displaced people, and in some regions, girls and young women—AI can be an academic and emotional lifeline. In Afghanistan, where the Taliban has banned women from post-primary education, girls and women have harnessed the potential of the Internet and AI to further their education. The School of Leadership Afghanistan has employed AI to digitize the Afghan curriculum, create lessons based on this curriculum, and disseminate content in Dari, Pashto, and English via WhatsApp lessons (Bottigliodivico 2025). More poignantly, faced with severe restrictions on their physical movements and social interactions, many Afghan girls and young women are using chatbots for emotional support and friendship (Clun 2025).

## **ADDRESSING CAPACITY CONSTRAINTS**

Low-resource education systems struggle with fundamental constraints—insufficient funding, lack of trained teachers, overcrowded facilities, and minimal time or resources for individualized support for learners. AI holds promise in addressing many of these constraints.

For example, AI technologies can address capacity constraints at the individual learner level by functioning as an independent, self-directed learning tool, potentially more affordable and accessible than traditional educational resources. Teachers interviewed for this study particularly recognize the potential of AI tools like LLMs and tutoring programs to benefit struggling learners, many of whom come from underresourced backgrounds, thereby expanding equitable access to educational support that was previously available only to those with greater resources (Burns 2021). LLMs can translate content into local languages while providing ubiquitous access to current information and accelerating the development of basic skills. Indeed, as one teacher from India notes, for students with high degrees of self-directedness and curiosity, LLMs are “a goldmine.”

At the systemic level, the economic accessibility and open-source nature of some AI tools further enhance their equity potential for schools. Free AI tools can reduce school software expenditures, while individual monthly subscription fees for more advanced LLMs, though prohibitive for many, remain affordable enough that several teachers in our study reported pooling resources to share a monthly subscription.

Additionally, AI technologies—including chatbots, LLMs, ITSs, productivity tools, and adaptive applications—can compensate for gaps in teacher knowledge while providing crucial scaffolding for paraprofessionals, emergency teachers, long-term substitutes, and educators teaching outside their area of certification. Low-cost AI tools, trial versions, and even free LLMs can be particularly valuable for critical teacher functions such as

lesson planning and assessment design, helping mitigate the impact of teacher absences and poor teacher quality by augmenting, not replacing, human teaching capacity (Burns 2023b).

Beyond their direct support to students, which will be discussed in Benefit 4, personalized AI programs also hold considerable promise for strengthening teacher capacity itself, addressing the instructional limitations that often constrain learning in underresourced settings. While research on student-facing applications dominates the literature, emerging evidence suggests that teacher-facing AI tools can compensate for deficiencies in content knowledge, provide essential scaffolding for para-teachers or unqualified instructors, and enhance overall instruction quality (Burns 2023b). Additionally, ITSs, grounded in predictive AI, can leverage AI to help teachers in low-resource contexts interpret and act upon the diagnostic data that ITSs provide, suggest instructional strategies, and draft follow-up lesson plans. This support is particularly valuable for teachers who may lack content expertise or pedagogical training, particularly in contexts where pre-service teacher preparation and ongoing teacher professional development remain unavailable or insufficient.

## **INDIVIDUALIZED ATTENTION IN CROWDED CLASSROOMS**

The personalization capabilities of AI (discussed in greater detail in Benefit 4) represent perhaps the most significant opportunity for addressing educational inequities, particularly in underresourced educational contexts globally. Research suggests that personalized learning, with and without AI, may help narrow learning gaps for marginalized students in the Global South, where class sizes are often large, undiagnosed learning issues common, and teacher capacity for individualized attention scarce. AI-driven personalized platforms allow teachers to create smaller instructional groupings working on differentiated tasks, freeing teachers to concentrate on students requiring individualized attention, and supporting tailored, developmentally appropriate

“teaching at the right level” approaches (Major and Francis 2020). Thus, personalized learning programs may narrow achievement gaps among marginalized students through high-dosage tutoring at significantly reduced costs, expanding access to educational support previously available only to those with greater resources (Burns 2021).

## **CONCLUSION**

While AI's benefits still favorably accrue to those with robust infrastructure, there are a number of technical innovations that can extend AI's benefits to underserved communities. Box 4 outlines several approaches to address these challenges in reach and affordability. Though issues around updating, handling diverse student queries, and ensuring device neutrality remain, such developments can potentially make AI access more equitable.

## BOX 4

# Possible solutions for expanding access to AI in underserved areas

**Small language models:** Small language models (SLMs) are language models with less computational complexity and a smaller amount of training. In contrast to LLMs, they consume less energy and require fewer resources, making them less expensive and more accessible for widespread use, particularly in low-resource settings (Kumar et al. 2025). SLMs trained on specific content or languages, when bundled with technologies students already possess and know how to use, like smartphones and messaging apps, could deliver low-cost, high-reach education. SLMs will be discussed in greater length under Risk 4.

**Open weight models:** These are models where the trained model parameters are freely available for use and distribution, eliminating subscription fees. DeepSeek, LLaMA and most recently OpenAI, have all released open weight models (OpenAI 2025b). Open weights refer to the final parameters—the weights and biases—of a trained neural network that determine how the model interprets inputs and generates outputs. When AI developers share these parameters, they enable others to fine-tune, adapt, or deploy the model for their own projects, though the weights themselves do not include the underlying training code or datasets (Open Source Initiative 2025).

These open-weight models can run on devices with 16 gigabytes of memory, operate offline for enhanced data privacy, and (in the case of OpenAI) are released under the permissive Apache 2.0 license, although licensing terms vary across different models (Huckins 2025). Open-weight models allow developers, researchers, and educators—especially in low-resource settings—to experiment with and build AI tools without needing access to proprietary application programming interfaces (APIs), costly cloud service fees, or restrictive licensing agreements—though they do not eliminate all costs.

**Offline functionality with device neutral access:** Mobile-based AI applications can mitigate local infrastructure constraints and reach users despite Internet or cellular network limitations (Stasenko et al. 2025). Trained AI models can be loaded onto devices, such as game consoles, “dumb phones,” through distribution to schools via Secure Digital cards, USB drives, or single-board computers like Raspberry Pi, thus enabling access to AI content—albeit with some degree of degraded functionality—without Internet connectivity. Platforms such as the World Possible’s Remote Area Community Hotspot for Education and Learning and Learning Equality’s Kolibri have established precedents for offline solutions in low-connectivity educational contexts (Learning Equality 2025; World Possible 2025).

As an example of the potential solutions above, Llama Filesan—an open-source project that simplifies running LLMs on personal computers—enables “local AI.” This is AI that runs on users’ devices, not in the cloud. This ensures privacy, offline accessibility, control over data, and (perhaps more importantly) reduced dependence on the dominance of “Big Tech” AI companies that provide LLMs (Stasenko et al. 2025, 46).

## Benefit 2: AI can optimize teacher time for greater focus on students

*AI is my favorite colleague.*  
— Teacher

Teaching, though often spoken of as a unitary activity, actually comprises a series of highly discrete tasks—planning, reporting, assessing, designing, instructing, and revising—with each containing an even more diverse array of subtasks. Understanding this complexity is essential when examining how AI is transforming educational practice.

Teachers report being enthusiastic AI users, particularly of LLMs, lesson plan generators, and student learning platforms. A separate internal, unpublished Brookings’ survey of 303 teachers in the U.S. and India revealed varying degrees of AI adoption among teachers with Internet access. Twenty-seven percent described themselves as “beginning to explore” AI, 37% reported “trying a few things,” 34% indicated “regularly integrating” AI into their practice, 24% reported creating specific AI activities, and 17% described training or supporting teacher colleagues in AI use. Only 3% reported non-use of AI. Since teachers could select multiple categories, these percentages reflect overlapping stages of adoption rather than discrete groups.

While these data are preliminary and limited in scope, the representation of different groups of teachers reporting non-use, exploration, experimentation, routine integration, and appropriation demonstrates patterns consistent with Hord et al.’s (2006) Levels of Use framework, which characterizes how teachers progress in adopting educational innovations, like AI, from initial

awareness through mechanical use to integration and refinement.

### AI AS A PRODUCTIVITY TOOL

At this relatively early stage of AI’s trajectory, one of its greatest benefits rests in its power as a productivity tool, typically used outside of classroom hours to enhance teaching and learning outcomes.

Teachers interviewed for this study use AI primarily for “task replacement”—automating activities traditionally performed by human beings (Quezzaire 2025). Analysis of coded data reveals that 78% of teacher responses indicated using AI primarily for productivity-related activities. This includes generating parent emails, grading and providing feedback, translating materials, creating worksheets, rubrics, quizzes, and lesson plans, developing content, tutoring students, creating individualized education plans, automating essay scoring, and checking work authenticity (Ferman et al. 2024). This pattern appears to be AI’s most dominant application and is consistent with data across multiple studies showing that teachers primarily use AI for productivity purposes (Celik et al. 2022; Quezzaire 2025; Roy et al. 2024; Deng et al. 2025; Code.org et al. 2025; Malek Ash 2025; Leiva et al. 2025).

The evidence shows that these productivity gains are both significant and measurable. A recent study by the Education Endowment Foundation demonstrated clear efficiency from the use of AI in education. During the summer term of 2024, a randomized controlled trial in the UK was conducted with 259 science teachers from Year 7 and Year 8 across 68 schools over a 10-week period. Teachers in the intervention group used ChatGPT to develop lessons and resources for their classes and were provided with an online implementation guide. In contrast, teachers in the control group were instructed to prepare lessons without using any AI tools (Roy et al. 2024).

The findings demonstrated that teachers using ChatGPT for lesson preparation experienced a 31% reduction in planning time, completing tasks in 56.2 minutes compared to 81.5 minutes for nonusers and thereby saving an average of 25.3 minutes per week. Crucially, this efficiency gain did not compromise quality, as an expert panel could not distinguish between resources created with and without ChatGPT assistance (Roy et al. 2024). These findings align with broader patterns observed in American educational contexts. A Gallup and Walton Family Foundation survey found that U.S. teachers who report using AI tools report even greater time savings—an average of 5.9 hours weekly, or 6 weeks over a school year of 37.4 weeks (Malek Ash 2025).

The impact of these productivity gains extends beyond individual time savings, suggesting meaningful implications for teacher workload and sustainable practice in education.

## THE AI TEACHING DIVIDEND

Task replacement through AI generates efficiency gains that benefit both teachers and students (Cambium Learning Group 2024). Teacher interviews confirm these findings, with educators reporting direct workload reduction as a result of AI use for productivity. Most critically, these efficiency gains offer a substantial “dividend” that can be reallocated and reinvested in more high-yield uses of teachers’ time that can directly benefit students (Impact Research 2024; Malek Ash 2025).

This optimization (the shift from low-impact to high-impact activities) occurs at two levels. First, teachers enhance decisionmaking processes and workflow, including curriculum sequencing for improved knowledge retention, resource allocation across programs, and analyzing student performance data to identify learning gaps and recommend instructional strategies (Celik et al. 2022). For instance, teachers in Singapore, one of our study sites, use AI to manage basic corrections like punctuation and capitalization, allowing both

teachers and students to focus on voice, style, and argumentation.

Second, the time saved from non-student-facing tasks can be reallocated to prioritizing direct engagement and relationship-building through progress monitoring, individualized feedback, targeted interventions, and personalized instruction (Celik et al. 2022). While more research is needed on how widespread these practices are and how teachers actually use this additional time, teacher interviews suggest that this optimization of teacher workflow may play a meaningful part in student learning outcomes.

## SUPPORTING TEACHER RETENTION

While task replacement offers significant benefits to individual teachers and students, it also holds potential value for the education system itself. This systemic impact becomes clear when examining the workload crisis facing the teaching profession.

With a few exceptions, educational systems across the globe are facing a dramatic teacher shortage. The United National Educational, Scientific, and Cultural Organization (UNESCO) estimates the need for 44 million additional teachers by 2030 in order to reach Sustainable Development Goal 4 of providing inclusive, equitable, and quality education for all (UNESCO and International Task Force on Teachers for Education 2030 2024). While numerous discrete and cumulative factors drive the global teacher shortage, workload plays a critical role. In the U.S., for example, teachers work substantially more hours than other professionals, averaging 53 hours per week during the school year, compared to 46 hours for non-teaching professionals and in England teachers average 48 hours of work per week (Doherty 2020; Steiner et al. 2023). Compounding this burden, as much as 25% of teachers’ work remains uncompensated (Steiner et al. 2023).

Evidence suggests that excessive workload is a key contributor to teacher attrition (teachers leaving the profession for reasons other than retirement)

and may in fact constitute “the crucial factor” influencing teachers’ decisions to leave teaching (Worth et al. 2018). Teachers experiencing burnout from the demands of teaching are more likely to leave the profession (Steiner et al. 2023). In England, “unmanageable workload”—marking, data collection, inspection requirements, paperwork, and frequent curriculum changes—represents the most frequently cited reason teachers give for leaving (Duff 2018).

Teacher attrition extends far beyond individual career decisions to threaten overall educational quality, particularly for learners most in need of high-quality educational interventions—those in refugee contexts, rural schools, urban schools serving children of color who are poor or migrants, and learners in the poorest global regions, including sub-Saharan Africa and Southwest Asia (Akiba et al. 2007; Ring and West 2015; Teacher Task Force 2021; United Nations Educational, Scientific and Cultural Organization 2022). Attrition is also enormously expensive for education systems, requiring schools and districts to shift financial and human resources away from other programs to recruit, hire, and train replacements (National Commission on Teaching and America’s Future 2007).

The Learning Policy Institute estimates that replacing one teacher in a large U.S. district costs an average of \$25,000 (2024). Thus, a U.S. school district with 5,000 teachers and a very conservative 12% attrition rate faces annual replacement costs of approximately \$15 million. Given such substantial financial implications, support like AI that reduce work demands and increase efficiency represent critical developments.

## CONCLUSION

While these time-saving benefits need to be combined with other reforms around pay and working conditions, anything that makes teaching more “efficient” offers welcome relief. There certainly are other innovations, such as reimagining the composition of the education workforce, that

can help address teacher shortages (Ingersoll and Audrain 2025). But alongside these, AI’s growing capacity to complete educator-related tasks, even complex ones, with greater speed and quality holds great promise for both teachers and students, offering expanded time savings and enhanced optimization of the essential human interactions at the core of teaching (Fisher 2025). Thus, provided that administrators protect this newly gained teacher time (see Box 5), the productivity and efficiency gains of AI may serve a supporting role in overall policies focused on addressing teacher workloads to enhance retention.

### BOX 5

## Time savings only occur if teacher workloads remain steady

The AI teaching dividend described here only materializes if education leaders maintain teachers’ current workloads. Participants in the study expressed concerns that time savings from using AI could be quickly eroded if administrators assigned additional duties. Some educators also worried about openly sharing the extent of their time savings, fearing that government leaders might view it as an opportunity to cut costs, do more with less, and significantly increase teacher workloads. Ultimately, understanding system dynamics is essential when considering potential outcomes of implementing innovation (Fuller and Kim 2022; Sengeh and Winthrop 2022).

## Benefit 3: AI can improve student learning

*AI explains things better.*

— Parent, student, teacher

AI provides a wide range of potential and actual benefits to students' learning and development through both direct and indirect pathways. As is often the case with new technologies, there is a lag between the use of AI and rigorous research documenting the accumulated benefits on student learning. This reality notwithstanding, interviews with study participants suggest that AI applications, when implemented in pedagogically and developmentally appropriate ways, can meaningfully enhance student learning outcomes across multiple domains. Research by and with education and technology companies also shows that embedding AI features—such as interactive question-and-answer opportunities—within textbooks and other static, text-based teaching and learning materials improves student learning outcomes across a range of subjects, from biology to history (Anderson 2025; LearnLM Team and Google 2025; Microsoft 2025a; Pearson 2025). In our study, 62% of teacher responses suggest that AI can improve student learning by empowering students. When asked which elements of learning will gain most from AI, 35% of teacher responses point to the two domains most traditionally associated with education: reading and writing, explored below.

### **SUPPORTING STUDENT READING, ESPECIALLY FOR SECOND LANGUAGE LEARNERS**

Long before the widespread adoption of AI, technology-enhanced learning has long provided a more engaging and effective approach to reading instruction and language instruction compared to traditional methods (Burns 2021). Teachers

interviewed for this study enumerated several elements of AI that lead to improved student reading comprehension—particularly for those learning a second language (L2 learners). Many of these advantages directly address the linguistic foundations of reading, including vocabulary knowledge and syntactic understanding. Examples include the capacity of AI to simplify vocabulary, modify reading comprehension questions, and adjust the Lexile (or reading difficulty) levels of texts to match individual learners' proficiency levels appropriately.

However, the benefits of AI also extend to important nonlinguistic factors that, while unrelated to language per se, remain crucial for reading comprehension. Teachers attest to the ability of AI to engage students in the reading process, scaffold learning strategies, extend reading time as needed, utilize text-to-speech features, and create zones of privacy and comfort where student reading difficulties remain confidential.

While research on AI and reading comprehension is scant at present, with studies of varying quality, the evidence that does exist supports teacher assertions regarding the capacity of AI to enhance reading comprehension for L2 learners. A range of studies, employing different methodologies, from a spectrum of countries (Kuwait, Saudi Arabia, India, and the Philippines), student age cohorts, and a range of AI tools (ITSs, personalized learning platforms, and general-purpose AI platforms) corroborate teacher claims (Alazemi 2024; Ali et al. 2023; Wangdi and Shimray 2025).

For example, a study from Kuwait appears to show promising results for AI's impact on reading comprehension and language-learning outcomes. Eighty grade 9 female English-language learners were selected through convenience sampling and divided equally into treatment and control groups comprising 40 students each. Researchers administered pretests that were conducted to assess their initial levels of English-language comprehension. The treatment group received AI-based instruction through a commercial

multimedia platform, *Nearpod*, while the control group received traditional face-to-face instruction. Both groups were taught 10 reading passages, and post-tests were administered to measure reading comprehension progress. While both groups exhibited parity in their pretest reading scores, the post-test showed a small improvement in the intervention group, which scored 17.75 compared to 16 for the control group. Statistical analysis confirmed that this 1.75-point difference was unlikely to have occurred by chance, indicating a genuine if modest effect (Alazemi 2024).

The above findings, along with information provided through teacher interviews, suggest that AI-based personalized reading platforms may be particularly effective in improving reading comprehension, especially for L2 learners. Much of this success appears linked to the ability of AI to match the level of reading difficulty to individual student reading abilities, analyze vast amounts of data to provide targeted insights, feedback, and personalized support to students, as well as to engage students in more individualized ways (such as asking more personalized questions), thus enhancing student engagement and motivation. Personalization will be discussed in the next benefit.

## HELPING STUDENTS WITH THE WRITING PROCESS

AI holds significant promise for enhancing writing instruction and student outcomes across multiple dimensions of the writing process when used to amplify student learning rather than shortcut it. Interviews with teachers and education experts, alongside emerging research, suggest that AI technologies can support both teachers and students across the entire writing continuum—when thoughtfully integrated into pedagogical frameworks. This integration requires balancing structured writing with more exploratory and iterative practices and maintaining equilibrium between human guidance and technological support (Ferman et al. 2025; Deng et al. 2025; Guo and Wang 2024; Dell’Acqua et al. 2023).

The use of AI tools to support student writing offers three benefits:

**Process-based writing support:** When used as part of a pedagogical approach that supports (not replaces) student writing, AI technologies can provide meaningful assistance in the writing process. Teachers in our study report that students are more motivated to write when using AI as a scaffolding tool. This effective and motivational support of AI can be particularly valuable for students who have writing anxiety, who are weak writers, or who have difficulty conceptualizing writing topics or formulating written thoughts. It can also supplement the traditional writing process in schools which often focuses on assigning, rather than teaching, writing.

AI can help students in all phases of the writing process. During the ideation and conceptualization phase, teachers report that AI can “spark creativity” and help students overcome writer’s block by helping them “get unstuck” so they can develop arguments and ideas. At the drafting stage, it can help with organization, coherence, syntax, semantics, and grammar. At the revision stage, AI can support the editing and rewriting of ideas as well as help with the mechanics of writing, including punctuation, capitalization, and grammar.

**Feedback:** In *Visible Learning*, his groundbreaking meta-analysis of over 50,000 education studies, John Hattie highlighted the profound impact of feedback on student learning outcomes, ranking it among the most powerful educational interventions in terms of average effect size (2010). Feedback plays a particularly important role in supporting students’ writing revision skills because it operates on both an academic and an affective level, helping students advance their writing skills while cultivating the confidence needed to complete writing tasks (Graham 2019; Meyer et al. 2024). Yet high-quality feedback is time consuming, and teachers can only support and guide a certain number of students at any one time. To counter this constraint, teachers are increasingly using technologies like LLMs, as well Automated Written

Evaluation (AWE) systems, such as Grammarly, to automate assessment and feedback. AWE systems typically comprise two essential components: a scoring engine that produces automatic scores, and a feedback engine that delivers Automated Written Corrective Feedback.

There is growing evidence that this array of AI tools can be used to identify individual writing needs and provide meaningful, personalized feedback to improve both the process and product of writing. Meyer et al. (2024) randomly assigned 200 German secondary school students to receive LLM feedback after writing an essay draft in English (the treatment group) while 250 students received no LLM feedback on their essay (the control group). Those receiving LLM feedback experienced very modest statistically significant improvements in text revision ( $d = 0.19$ ). However, more notable benefits occurred in how students felt about writing (motivation and emotions) with medium improvements in task motivation ( $d = 0.36$ ), and positive emotions ( $d = 0.34$ ).

Though these impacts are quite modest, the researchers advocate for the use of AI to create feedback opportunities for students that would be otherwise unavailable given teacher workload: “Even if the beneficial effect of one single LLM-generated feedback session is relatively small for revision performance, the effects of further feedback could cumulate for students because the automation allows for multiple feedback cycles and opportunities for practice” (Meyer et al. 2024, 7).

While AI demonstrates strengths in writing support, research also emphasizes the complementary nature of AI and human feedback. The University of California at Davis’s curricular intervention, Peer & AI Review + Reflection (PAIRR), blends both AI and human feedback to support students during the writing process, following up with critical reflection on the process and its impact on their learning (University of California at Davis 2025). One study of 400 feedback instances for the same essays (200 pieces of human-generated formative feedback and 200 pieces

of AI-generated formative feedback) found that human raters provided superior feedback in terms of clarity, accuracy, supportive tone, and emphasis on critical improvement areas, while AI feedback excelled in delivering criteria-based evaluations (Steiss et al. 2024). This research suggests that optimal writing instruction may benefit from hybrid vigor, leveraging and integrating both AI and human feedback with AI (in formative early drafts or instances where a well-trained teacher is unavailable) but relying on human teachers for more advanced stages of the writing process.

**Improved writing outcomes:** A number of quantitative studies (Meyer et al. 2024; Guo and Wang 2024) and qualitative ones (Meniado et al. 2024; Wang 2025) point to the role of AI, when used as “an active collaborator,” to create demonstrable improvements in student writing (Meniado et al. 2024, 7). As with AI tools for reading, these benefits are especially pronounced for L2 learners and involve the production of coherent texts with clear ideas, improved writing flow, and discourse maturity that approaches native-speaker-like quality, with significantly higher indexes of lexical diversity reflected in both ideation and sentence fluency. AI does this by accelerating the writing process, easing cognitive load, fostering new learning opportunities, providing immediate feedback, and promoting positive feelings about writing (Wang 2025).

While AI for writing studies examined in this study focuses on LLMs, AI-based writing tools have become increasingly popular writing supports. Kaliisa et al.’s (2025) meta-analysis combined findings from 41 separate studies involving nearly 5,000 secondary and post-secondary students to examine the effectiveness of AI feedback on student writing. The research determined that students who received AI feedback learned just as much as those who received traditional teacher feedback, suggesting that AI feedback is essentially equivalent to human feedback in terms of measurable learning outcomes. Rather than advocating for one approach over the other, the researchers recommend that educators use both

strategically. This hybrid model leverages what AI does well—providing quick, consistent feedback at scale—while preserving what human teachers do best—offering empathetic, nuanced, and deeply contextual responses.

An important caveat in terms of examining the effects of AI to help student writers focuses on the type of AI used. Many studies examining AI tools in education conflate fundamentally different types of AI systems, such as rule-based approaches, machine learning models, and generative AI. This conflation is problematic because these technologies operate on distinct principles and offer qualitatively different capabilities.<sup>1</sup> Treating these diverse technologies as equivalent when evaluating their educational impact obscures important distinctions in how they function and what they can reasonably accomplish in supporting student learning.

## CONCLUSION

Given the prevalence of reading and writing difficulties among students, especially L2 learners, the above research findings hold important implications for teachers seeking to enhance students' reading skills. The unique linguistic capabilities of AI tools such as LLMs can extend support for struggling readers and writers who require help in decoding, comprehending, and communicating in the conventions and syntax of a non-native language (Marzuki et al. 2023).

Helping struggling readers and writers is often a major focus of education systems everywhere. Many schools invest heavily in specialized literacy intervention programs, additional textbooks, outside specialists, and remedial software. Reading and writing assistance is often highly labor intensive. Though far more research is needed, AI tools may potentially reduce the intensive one-on-one human support traditionally required for struggling readers and writers, particularly L2 learners or those with disabilities, and might serve as an integral element of structured literacy interventions. From a human resource and financial perspective, these

developments could alleviate financial and human resource shortages in education systems struggling with both.

## Benefit 4: AI can tailor learning to each student's needs

Education is often maligned for its “one size fits all approach” to teaching and learning. For study participants, AI is changing that paradigm through two primary mechanisms: personalized learning programs and ITSs, which they rank as an important benefit of AI.

### PERSONALIZED LEARNING PROGRAMS

Personalized learning tailors instruction to each student's strengths, needs, and interests while providing choice, voice, and flexibility in reaching learning outcomes (Slocum 2016). Although personalization doesn't require technology, the logistical burdens of addressing every student's unique needs means that teachers increasingly *employ* technology, such as AI-powered platforms using LLMs or chatbots to adapt curricula, pacing, and learning environments to individual student profiles.

Teachers in the U.S., Nigeria, India, Singapore, and South Korea interviewed for this study report leveraging AI primarily for strategic adaptations of materials, instruction, and assessments; the tools provide immediate, responsive feedback. These practitioner observations align with research demonstrating that AI technologies facilitate adaptive pacing by responding to individual trajectories, enable continuous assessment with real-time adjustments, enhance engagement through dynamic experiences, and expand opportunities for student questioning (Niemi et al. 2023; Ferman et al. 2021).

Survey data from Cambium Learning Group (2024) confirms this trend: 56% of U.S. teachers currently leverage AI for personalized experiences, 52% employ these tools for real-time performance tracking and feedback, and 50% use AI-powered platforms specifically for developing critical thinking skills, suggesting that personalization extends beyond basic remediation to encompass higher-order cognitive development.

As discussed in Benefit 1, personalized learning can meet diverse learning needs of large student populations with constrained educational resources while offering high-dosage tutoring at low costs (Burns 2021, 37). Personalized systems do this by providing just-in-time feedback, recommendations, explanations, and examples, thereby improving students' conceptual development (Ng et al. 2024, 1333).

Major, Francis, and Tsapali's 2021 meta-analysis of 16 randomized controlled trials across five low- to middle-income countries involving 53,029 learners aged 6–15 years found that technology-supported personalized learning had a statistically significant positive effect size of 0.18 on learning outcomes ( $p = 0.001$ ).

Critically, their meta-regression revealed that adaptive approaches, which adjust to learners' levels, produced significantly greater impact (an effect size of 0.35, considered moderate to strong) than those only linking to learners' interests or providing personalized feedback and assessment (Major et al. 2021, 1935). This demonstrates that adaptive personalization—where instructional content, pacing, or difficulty dynamically adjusts based on current knowledge level or ability—produces substantially stronger learning outcomes than other forms of personalization.

## **INTELLIGENT TUTORING SYSTEMS: EVOLUTION AND EVIDENCE**

In the 1980s, Benjamin Bloom identified education's "2 sigma problem"—namely, that students receiving one-to-one tutoring scored two standard deviations

higher than those without tutoring. The "problem" was making such intensive support scalable and economically feasible (Bloom 1984). Over the past several decades, technology, in the form of computer-based and online tutoring, has been harnessed to provide a scalable, cost-effective alternative to in-person tutoring.

ITSs represent the most common AI educational applications, with millions of users worldwide (Burns 2021). These systems have evolved from early rule-based expert systems using predictive algorithms and knowledge tracing to sophisticated platforms incorporating AI for natural dialogue, dynamic problem generation, and affective data analysis (du Boulay 2016; Buckley et al. 2021). ITSs typically employ cognitive models grounded in learning science theory, using techniques like model tracing or constraint-based modeling to diagnose student knowledge states and provide targeted instructional interventions.

Research consistently demonstrates their effectiveness. Kulik and Fletcher's (2016) meta-review of 50 studies concluded ITSs can "match the success" of human tutoring, while Nickow et al. (2020) found an average effect size of 0.37 standard deviations across 96 randomized controlled trials—equivalent to moving students from the 50th to 65th percentile. Benefits prove strongest in foundational subjects and are particularly helpful for traditionally underserved populations, including learners with disabilities (Kasneci et al. 2023).

While recent ITSs have begun incorporating machine learning methods—including neural networks and natural language processing—to improve student dialogue and predict learning trajectories, most deployed systems with established research evidence still depend on structured, rule-based cognitive architectures from the 1980s and 1990s (García-Méndez 2025). The introduction of LLMs marks a significant new phase in computer-based tutoring, expanding possibilities for flexible conversations and content generation. Granted, research supporting large-scale,

classroom-tested implementations of AI tutoring programs remains limited, lagging behind traditional ITSs that have demonstrated effectiveness through decades of rigorous evaluation (Nye et al. 2023, 2; Jiang et al. 2025). However, three recent studies offer early evidence of the considerable promise that AI-enhanced tutoring systems may hold.

In the first, researchers conducted a randomized controlled trial with 194 undergraduate physics students at Harvard University, comparing learning outcomes and student perceptions between an AI tutor and active learning classroom instruction (Kestin et al. 2025). Both instructional approaches were designed using identical pedagogical best practices. Students alternated between conditions across two weekly lessons on surface tension and fluid flow, completing pre- and post-tests to measure content mastery alongside surveys assessing their learning experience.

Students using the AI tutor demonstrated significantly higher learning gains in less time than those in active learning classrooms. The AI group achieved a median post-test score of 4.5 compared to 3.5 for the classroom group—representing more than double the learning gains relative to the combined pre-test baseline of 2.75. Students also reported feeling more engaged and motivated when working with the AI tutor. The authors conclude that students interacting with the AI tutor at home “learn significantly more than when they engage with the same content during an in-class active learning lesson, while spending less time on task,” suggesting substantial potential for AI tutors in authentic educational settings (Kestin et al. 2025, 4).

The second study, a 2024 World Bank trial in Nigeria found that Microsoft Copilot (generative pre-trained transform, version 4 [GPT-4]) improved first-year secondary students’ English skills by 0.23 standard deviations and digital literacy by 0.31 standard deviations—learning gains equivalent to 1.5 to 2 years of regular schooling in just six weeks. Researchers noted that the program succeeded because it complemented the efforts of teachers

rather than replacing them (Di Simone et al. 2025). Finally, Stanford’s Tutor CoPilot, which provides AI-guided support to human tutors, increased student mastery rates by 4 percentage points overall and 9 percentage points for students with lower-rated human tutors (Wang et al. 2025). At a cost of only \$20 per tutor annually, this may represent an important step in solving Bloom’s Sigma 2 problem.

## **CONCLUSION: THE IMPORTANCE OF “HYBRID VIGOR”**

Despite promising outcomes, AI-powered personalization raises substantial concerns among our study participants. The primary apprehension centers on equity: personalization may also function as sophisticated tracking, placing students on limited pathways where basic skills are reinforced to the detriment of more important competencies. AI’s predictive analytics risk creating self-fulfilling prophecies, whereby algorithmic predictions based on early performance or demographic factors shape educational opportunities in ways that confirm biases. Several Delphi panelists expressed worry that the poorest education systems will employ AI to teach students while students in wealthy education systems will be taught by AI—and human teachers.

Design limitations compound these worries. Because of ease of design, many AI-driven tools employ outdated theories of learning, such as behaviorism, thus limiting both what and how students learn (Chen et al. 2025; Bali 2024). While AI can produce vast quantities of materials leveled to a student’s learning ability, it often lacks the nuanced judgment needed for complex tasks like teaching reading comprehension or writing—where, in the words of one teacher, “human insight is irreplaceable.”

The real danger lies in using tutoring and personalization systems for substitution versus supplementation. When systems position learners as isolated individuals interacting with platforms, they *automate* rather than *augment*

learning, eliminating the embodied experiences and relationships so critical to human learning (Cukurova 2024). AI cannot discern whether student slowness indicates deep thinking, peer collaboration, fatigue, or disengagement—a contextual understanding that good educators develop through daily interaction.

Furthermore, the relational and affective dimensions crucial to genuine personalization—formative feedback, motivational scaffolding, shared goal setting—remain beyond algorithmic capability. Students derive motivation from interpersonal connections; if teachers become sidelined in favor of AI tools, in the words of one Delphi panelist, students may feel devalued as “algorithmic data points rather than individuals whom teachers genuinely care about.”

The evidence suggests AI-based personalization and tutoring work most effectively through human-AI collaboration. One teacher explains what this looks like in practice: “The students come back to me in the small group rotation. The personalized platform has helped them with capitalization, punctuation, and grammar. Now I can help them extend their thinking further.”

This “hybrid vigor” ensures that technology augments, rather than replaces, the human dimensions of meaningful learning. At its best, AI-powered personalization helps each student feel their learning is meaningful while providing teachers with feedback on student progress to enable tailored instruction. The most sophisticated AI-powered methods require a strong teacher-student dialogue to succeed, maintaining the relational foundation essential to education while leveraging technology’s capacity for scale and individualization.

## **Benefit 5: AI can extend learning to neurodivergent students and students with disabilities**

Education systems often struggle to meet the unique needs and strengths of all students. When schools are not accessible and inclusive, participation and learning become especially challenging for students with disabilities and neurodivergent learners. Children with disabilities are nearly 50% more likely to have never attended school. Of the 240 million children with disabilities worldwide, more than 90% in low- and middle-income countries lack access to the assistive technology they need (United Nations Sustainable Development Group 2024). Since most schools are designed for neurotypical children, neurodivergent students—such as those with dyslexia, attention deficit hyperactivity disorder, and autism—are often left out (Teach Access 2025).

AI can help education systems better support children with special needs by expanding the capacity and availability of assistive technology. Globally, more than 2.5 billion people, including both adults and children, need assistive devices such as wheelchairs, hearing aids, or apps that support communication and cognition. Yet nearly one billion lack access, particularly in low- and middle-income countries where availability meets as little as 3% of need (World Health Organization and UNICEF 2022). As we examine next, AI offers promising solutions to address this accessibility gap through applications ranging from adaptive assessments to innovative delivery mechanisms.

## SUPPORTING NEURODIVERGENT STUDENTS

Since the development of the first chatbot—ELIZA—at the Massachusetts Institute of Technology (MIT) in the 1960s, advances in machine learning and natural language processing have produced increasingly sophisticated conversational agents. Social chatbots or artificial companions, including companion and emotional chatbots, which will be discussed in Section IV, engage in empathetic conversations with humans, deliberately incorporating anthropomorphic features to promote the development of social skills. These design elements strategically cultivate emotional connections between users and technology, creating more intuitive and emotionally resonant interactions (Chaturvedi et al. 2023).

Interviews with teachers and research data suggest these applications particularly appeal to individuals with autism spectrum disorder (ASD). Youth with ASD often experience significantly more social difficulties and maintain fewer friendships compared to typically-developed peers. Learners with ASD may gravitate toward online chatbots because they simulate real-life communication scenarios, providing a safe space for rehearsing interactions without risks of judgment or rejection (Franze et al. 2023; Ali et al. 2023; Zhao et al. 2025). One teacher in England illustrates these findings with an example from her classroom:

*I've got a lot of neurodivergent children. And they struggle with timing, pacing, and particularly structure. Once they've been shown like, okay, if you use these prompts, you can get an AI to help you plan or structure something, they find that very helpful. And they like the fact that they can ask as many questions as they want.*

Technology can also be helpful in supporting students with dyslexia. Approximately 15–20% of the global population has dyslexia and it is often underdiagnosed (International Dyslexia Association, n.d.). Without relevant support, students

with dyslexia struggle to learn in text-heavy environments like schools because they struggle with things such as spelling, letter order confusion, slow reading, handwriting, and phonological processing (Kuerban et al. 2025).

For example, Let AI Read First (LARF) is an AI-annotated text approach designed to enhance reading abilities of individuals with dyslexia. The system uses preset prompts to guide GPT-4 in processing original text and incorporating HTML tags to manipulate text display through bolding, highlighting, italics, font color changes, and size adjustments, while textual content remains entirely unchanged.

Zhao et al. (2025) conducted a between-subjects experiment with 150 participants who self-reported having dyslexia or suspected undiagnosed dyslexia, all with English as their mother tongue. Participants were randomly assigned to three conditions: Control (unmodified text), Conventional (Bionic Reading, which bolds the first few characters of words), and LARF (GPT-4 annotated text). They were assessed on reading functions including comprehension, accuracy, and retrieval. LARF was helpful for participants with severe dyslexia, with improvements in reading performance such as retrieval, recall and comprehension being more pronounced in those experiencing more significant challenges.

Other approaches integrate augmented reality (AR) and AI. ReadSmart, for example, is an application that captures text from real-world environments using AR devices, processes text using optical character recognition (OCR), summarizes and generates educational content with AI, and displays this content in an interactive AR environment (Kuerban et al. 2025, 164). AR is a semi-immersive interface that allows a combination of real-world elements captured through a camera with multimedia elements such as text, images, and video (Burns 2023b).

Researchers examined how AR and AI could be integrated to design an educational system that

effectively supports students with dyslexia. They subjected the platform to extensive performance testing across 100 trials. User testing via surveys included 45 professors/teachers and 258 students in the first iteration, and 40 participants including 4 professors/teachers, 25 students, and 13 professionals in the second iteration. The platform received an overall satisfaction rating of 3.9 out of 5, with particularly positive reception for text-to-speech and OCR accuracy (Kuerban et al. 2025).

Both examples point to broader implications for integrating AI with assistive technologies to create multimodal learning environments that address the specific challenges faced by students with dyslexia.

## **SUPPORTING STUDENTS WITH DISABILITIES**

Many learners with disabilities rely on assistive technologies—hardware, software, digital content, apps, and peripherals modified to assist individuals with disabilities. Different assistive technologies serve various support functions: adapted trackballs and alternate keyboards help students with motor impairments; screen readers like JAWS (Job Access With Speech), digital talking books, and voice-to-text software assist visually impaired students; communication apps, speech synthesizers, and sign-language resources help students unable to use verbal speech.

Most powerful perhaps may be the use of AI for those with speech impairments. The human voice has often been referenced as our “acoustic fingerprint”—a unique human identifier (Costello 2000). Advances in AI combined with innovative assistive technologies can help both children and adult learners reclaim or preserve their voices.

For many adults living with conditions such as aphasia, ALS or laryngeal cancer, voice-cloning technologies can create synthetic replicas of their original voice, increasingly in ways that are indistinguishable from natural voices. Voice-banking techniques—where a user records their voice in advance for subsequent use—can create

personalized synthetic voices that can be used with communication devices (John Costello, personal communication, November 12, 2025).

For children with speech impairments, voice-banking and emerging synthetic-voice systems can provide a personalized synthetic voice that preserves elements of the child’s vocal character. They can be used with Augmentative and Alternative Communication (AAC) devices as children grow and their needs evolve (though the banked or synthesized voice cannot be aged as the child matures). AAC is another major recipient of AI-enhanced access. Some communication apps now integrate AI features such as predictive text to reduce keystrokes, or speech-recognition modules tailored to non-standard speech patterns arising from motor-speech disorders, thereby helping users who cannot speak verbally to communicate through synthesized voice output (Teach Access 2025).

Beyond voice restoration and AAC, AI-powered tools are broadening communication access in multiple ways: advanced text-to-speech and speech-to-text systems with predictive typing are aiding written communication; voice assistants support visually-impaired learners in navigating digital libraries; augmented-reality apps are in development to translate printed text into sign-language for deaf students; and speech-enabled, socially assistive robots are being piloted to provide structured interaction opportunities for learners with ASD (Miao et al. 2021, Burns 2021).

## **CONCLUSION**

Finally, and critically, for instructional designers creating accessible online learning experiences, AI provides substantial support: generating alternative text for images, performing automated accessibility auditing using readability metrics, automating contrast checking for colorblind learners, verifying appropriate fonts for visually impaired students, and generating closed captions for video content. One small field experiment found that when novice designers alternated between using AI assistance and working independently, AI-assisted lessons

received significantly higher quality scores on half the assignments and never scored lower on average, suggesting AI's potential to elevate both accessibility and overall quality of online course design (Moore et al. 2025).

## Benefit 6: AI can advance assessment

*AI can support the analysis of multimodal portfolios, texts, audio, video, images, and offer insights that help teachers provide more meaningful feedback.*

— Delphi panelist

*We can use [virtual reality] and AI to assess the capacity to do, which is ultimately more helpful and valuable in the working world...aviation has been doing [this] for decades in flight simulations.*

— Delphi panelist

*And not only did [the AI assessment platform] mark [the exam] for her, it then highlighted the areas that she needed to work on...From a student's point of view, it's honing in on her learning [and] what she still needs to work on. I was kind of blown away with that.*

— Teacher

As the quotations above demonstrate, assessment is both critical to and inseparable from teaching and learning. AI's potential to transform assessment practices is one of its most significant contributions to education. While many teachers interviewed for this study shared their struggles with AI and assessment around issues of plagiarism and cheating—consequently adopting AI-safe assessments, such as oral or paper-and-pencil exams—other teachers have embraced AI's potential for assessments that are more accurate and accessible, as well as capable of tracking a wide range of competencies. Three dimensions

illustrate how AI is potentially transforming assessment: automation, effectiveness, and expansiveness.

### AUTOMATION

For teachers interviewed for this study, AI has significantly advanced the automation of assessment across multiple dimensions, including exam construction, test item generation, automatic grading, and the recording and reporting of data at scale. This automation addresses a persistent challenge that many national assessment systems have faced: the difficulty of gathering sufficient data to draw reliable conclusions without creating undue burdens on teachers, imposing excessive costs on education systems, or delaying results that learners need in a timely manner (Burns 2023b).

The automation of assessments through AI offers three particularly valuable educational benefits. First, it lowers the cost differential of testing, because automated processes require significantly less time to score and store assessment data. Second, it enables a quick turnaround of assessment information to teachers and students, which allows educators to assess student performance at a much more granular and detailed level than was previously feasible. Third, automated systems can provide more reliable scoring and valid data interpretation, reducing the variability and potential bias that can occur with human scoring (Burns 2023b). Together, these capabilities represent one of the most significant contributions AI makes to education: the ability to reduce manual labor so that essential tasks are performed more efficiently, at higher volumes, more objectively, and at a scale that would be impossible to achieve through traditional methods alone.

### EFFECTIVENESS AND EXPANSIVENESS

Until AI, technology's greatest benefit for assessment was in terms of efficiency—automating and scaling the delivery of assessment and reducing the labor and cost of designing assessments. Arguably, technology

did less to address effectiveness, particularly good assessment design, which has long been challenging for teachers creating their own assessments. Developing good assessment questions, particularly those that assess not simply declarative and procedural knowledge but conceptual, higher-level thinking skills, aptitudes, and a full range of cognitive and affective skills has typically required highly skilled content experts, psychometricians, designers, and assessment experts.

AI is changing assessment in fundamental ways. While AI tools like LLMs and teacher productivity platforms can streamline the creation of valid, reliable, and transparent assessments, their more transformative potential lies in expanding what assessment can measure and how. Evidence from studies, in addition to discussions with teachers and Delphi panelists, suggests that AI enables entirely new forms of formative, process-based, multimodal assessments that analyze students' work as they progress through learning tasks, tracking a wider range of competencies than traditional approaches allow (Curriculum Associates 2025). This emerging paradigm supports more diverse pedagogical approaches and can accomplish this expansion in several ways.

First, through its ability to diversify assessment, AI makes it possible to administer multiple tests and measure student learning in multiple formats—via quizzes, discussions, projects, performance-based assessment, and e-portfolios—thus assessing knowledge and facts as well as higher-level skills, such as students' ability to analyze, synthesize, and evaluate ideas, arguments, and information. In this way, it becomes possible to assess a student's learning progress as well as their process and product (Delphi panelists). The administration of multimodal forms of assessment can capture a wider range of indicators from multiple data sources to assess the strengths and weaknesses of each student's approach (Blikstein and Worsley 2016; Di Mitri et al. 2018; Chango et al. 2021; Delphi panelists). These capabilities enable AI to provide personalized feedback, facilitate student

self-reflection through interactive dashboards, illuminate collaboration patterns and socioemotional dynamics typically invisible to educators, and offer real-time guidance to both students and teachers (Deng et al. 2025; Delphi panelists).

Second, adaptive assessment addresses the reality that students in any classroom occupy different points along the learning trajectory. These assessments automatically adjust difficulty based on student performance, administering easier tasks to those who struggle and more challenging ones to those who succeed. Unlike traditional assessments that report results against common content standards, adaptive assessments measure students against a learning continuum, estimating where each individual has reached regardless of age or grade level (Masters 2021). While adaptive assessments have long preceded generative AI, AI deepens their utility by selecting the most appropriate next question for each student, more accurately estimating ability by drawing on diverse datasets that handle uncertainty and detect anomalies, and providing real-time, actionable feedback that helps teachers and learners identify strengths, address misconceptions, and plan next steps (Ihichr et al. 2024; Cobing 2025).

Next, the natural language processing capabilities of AI, particularly LLMs, can enable more flexible, ongoing, and interactive assessment methods (Delphi panelist). For example, AI makes possible the use of longer free-text or audio responses, facilitating in-depth assessment of students' knowledge, mental models, and reasoning abilities beyond what conventional non-AI assessments typically achieve. Teachers can engage students in extended question-and-answer sessions where AI analyzes responses to identify areas of proficiency and focuses on demonstrated knowledge gaps or ambiguous understanding. The flexibility inherent in AI systems allows for sophisticated pattern recognition that can identify, assess, and address common student misconceptions through detailed analysis of student responses and close examination of language use. Thus, the natural language processing capabilities of AI may facilitate

more responsive and accurate assessments that adapt to individual student needs while maintaining rigorous evaluation standards (Delphi panelists).

Finally, combined with tools such as VR, AI can potentially assess students' capacity to perform in realistic contexts, which may help educators better connect students' knowledge acquisition to their knowledge application. Immersive environments, in combination with AI, can assess competencies ranging from critical thinking and problem-solving to communication skills within contextualized scenarios (Burns 2023b). They can enable iterative practice in, for example, practicing job interviews with AI-enabled avatars that respond in real time, creating consistently high-quality learning experiences previously available only to students with access to expert mentors and shifting from assessment practices prioritizing precision and ranking to complex, project-based assessments measuring agency and critical thinking across authentic scenarios. This VR-AI integration can allow educators to assess not merely what students know but what they can actually do in professional contexts, moving beyond information transfer to practical competency demonstration (Delphi panelist).

## **SUPPORT THROUGH PREDICTION**

Beyond classroom assessment, AI's capacity to identify patterns in vast datasets enables early identification of students who need support before challenges become crises. Researchers from NOVA University Lisbon analyzed the academic achievement of all 110,627 public secondary school students in Portugal in the 2014–2015 academic year to predict those at risk of failing. Researchers compared multiple AI methods (for example, random forests, artificial neural networks, etc.) against traditional logistic regression, using data from the beginning of each academic year to predict outcomes at year's end (Cruz-Jesus et al. 2020).

The AI methods demonstrated superior performance. Random forests correctly identified

the 5% of students most likely to fail with the highest accuracy (87%), while artificial neural networks correctly classified 80% of at-risk students compared to 49% for logistic regression. Researchers noted that these results “clearly demonstrate that AI approaches manifestly outperform traditional methods” for identifying students at risk of dropping out (Cruz-Jesus et al. 2020, 7). However, they cautioned that AI's technical capacity must be balanced against pedagogical wisdom, ethical considerations around educational equity, and the reality that prediction is not destiny.

## **CONCLUSION**

Identifying at-risk students represents only the first step—the critical challenges lie in determining how to help them.

This is where generative AI, a fundamentally different technology from the AI models used in the Portugal study, could play a transformative role. By synthesizing each student's risk profile with contextual information and evidence-based practices, generative AI could translate these statistical predictions into actionable interventions: personalized support recommendations, individualized learning plans, family communication materials, and specific resources tailored to each student's circumstances. This integration could transform early warning systems from diagnostic tools that identify problems into comprehensive support frameworks that guide educators toward targeted solutions, helping ensure that early identification leads to timely, appropriate intervention.



## IV POTENTIAL OUTCOME 2: AI can diminish learning if overused with no guardrails

As noted, AI's educational benefits are manifold. For many teachers interviewed, AI's capacity to save time, simplify tasks, and optimize work processes has freed them to provide more individualized student support, feedback, and tutoring. They report that AI has reduced bias and increased objectivity in grading, enabled the creation of differentiated content based on student interest and reading levels, and helped teachers interact more effectively with students who learn in nontraditional ways.

AI, particularly the LLMs most prevalent in the educational contexts examined in this study, is an extraordinarily powerful and versatile set of productivity tools that can effortlessly generate diverse artifacts and outputs. Yet this very source of AI's strength also constitutes its greatest risk. This tension emerged clearly in research findings. For study participants—experts, teachers, parents and students—the risks of AI far outweigh its benefits: 56% of all responses point to AI's harms versus 44% to its benefits.

Figure 7 displays the proportion of “harms” or “risks” of AI identified by study participants: experts, teachers, students and parents.

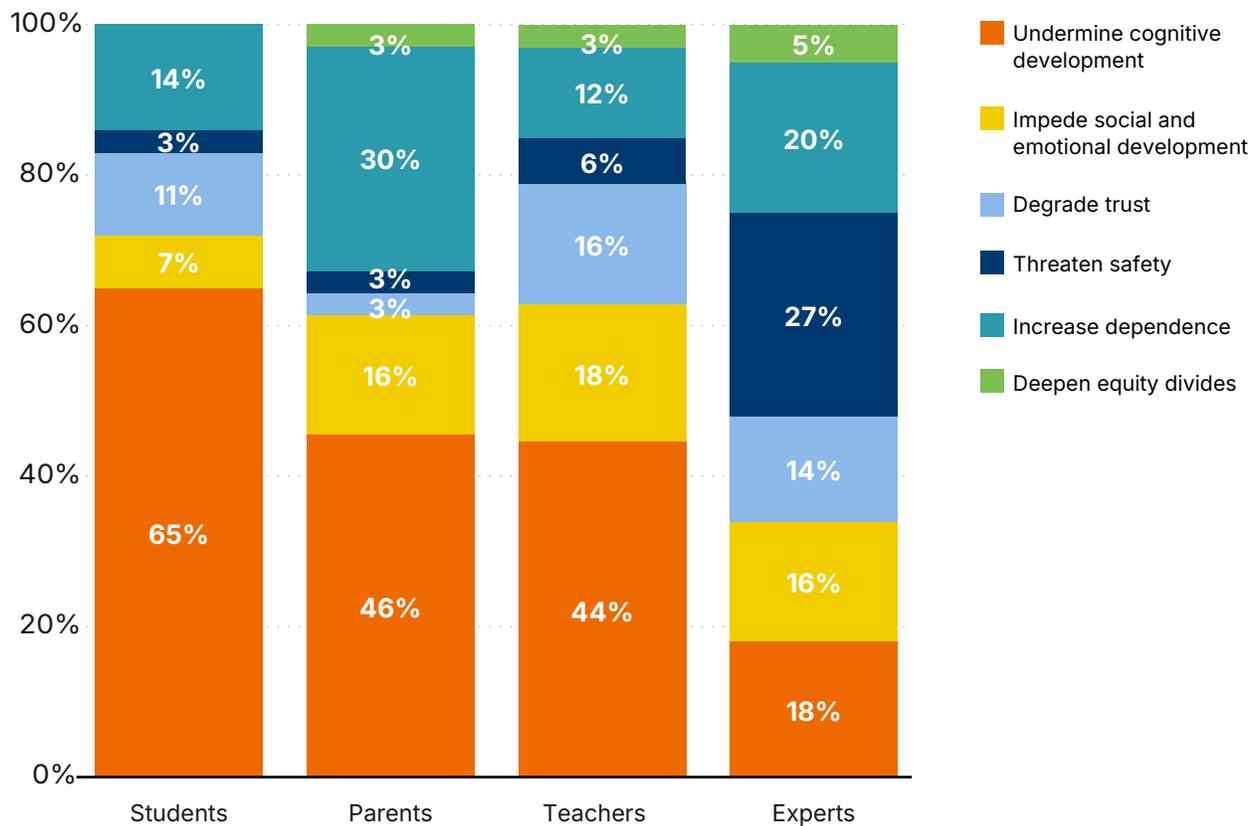
As Figure 7 illustrates, AI poses risks across multiple dimensions. These include developmental risks such as undermining students’ cognitive development, which includes content-based skills, critical thinking, and durable or transversal skills; social-emotional risks including dependency, isolation, and mental health impacts; privacy and data risks involving surveillance and security breaches; risks related to misinformation, bias, and inaccuracy in AI outputs; relational risks that erode student-teacher trust and replace human

interaction; and academic integrity concerns around authenticity and plagiarism. Further, AI’s benefits accrue mainly to students in well-resourced countries or education systems, thus exacerbating existing inequities.

The risks result from a combination of *internal* and *external* factors. Internal factors relate to the design of many AI tools and platforms, which encourage attachment and prolonged engagement while failing to provide cognitive scaffolds, adequate guardrails, or age-appropriate protections.

**FIGURE 7**

**Risks identified by study participants**



In this figure: teachers n = 117; experts = 72; students = 68; parents = 41.<sup>2</sup>  
 © Center for Universal Education at the Brookings Institution, 2026.

## BOX 6

# AI and early childhood development

While this study focuses on school-aged children, it is important to recognize the potential risks of integrating AI into interactions with very young children. As AI becomes embedded in toys and companies experiment with LLMs for early childhood, child development experts urge caution (Roche et al. 2025). As Dana Suskind, professor at the University of Chicago reminds us: “relationship maps laid down early have implications for how you interact with other humans for the rest of your life,” (Dana Suskind, personal communication, September 15, 2025). AI could influence the attachment process between infants or young children and their parents or caregivers—a foundation for healthy development and learning. Experts in our study are particularly concerned about vulnerable children with limited access to quality childcare and the temptation to use AI as a substitute for human interaction. Far more research and policy attention is needed to understand how AI may affect child development, especially for children from birth to age three.

External factors stem from the context and manner in which students use AI. Here the potential for harm is most acute when AI is used with minimal adult supervision and guidance, when AI replaces rather than enhances thinking and human interaction, and when students lack adequate preparation for productive AI use. These circumstances reflect the current reality for many students who encounter AI across multiple platforms throughout their daily lives—often outside school settings and increasingly beyond the reach of teachers and parents who might help them navigate these tools thoughtfully. Compounding these contextual risks are pedagogical approaches that encourage AI substitution rather than enhancement, issues of student motivation and maturity, concerns about academic integrity, and the absence of clear guidance and protection around appropriate and safe use.

We turn now to a discussion of these risks.

## Risk 1: AI can undermine students' cognitive development

*Students can't reason. They can't think. They can't solve problems.*

— Teacher

*If students can just replace their actual learning and their ability to communicate what they know with something that's produced outside of them and get credit for it, what purpose do they have to actually learn?*

— Teacher

*AI is the fast food of education.*

— Expert

*It's easy. You don't need to (use) your brain.*

— Student

We have long known about technology's impact on brain physiology (Carr 2010). In *The Anxious*

*Generation* (2024), Jonathan Haidt points to “great rewiring” of the brains of children and youth because of social media. AI continues this pattern.

The data presented here points to what might be termed a “great unwiring” of students’ cognitive capacities. The overuse, and even the routine use, of AI is fundamentally and negatively reshaping how students approach learning; this has profound implications for cognitive engagement, the development of practical skills, and existential questions about human purpose and agency. Study participants identified this threat to student learning as the greatest risk posed by AI. Fifty-seven percent of all responses discussing potential harms focused on threats to students’ “cognitive development.” This concern resonates across cohorts, appearing in 65% of student responses, 46% of parent responses, and 44% of teacher responses addressing AI’s potential harms.

---

### The data presented here points to what might be termed a “great unwiring” of students’ cognitive capacities.

---

Cognitive development is defined here as “the growth and maturation of thinking processes of all kinds, including perceiving, remembering, concept formation, problem solving, imagining, and reasoning” (American Psychological Association 2018).

The concern among study participants is that the routine use and overuse of AI do not simply harm student’s cognitive development—both actively place children at risk of cognitive decline.

The risk stems from what study participants view as a straightforward progression: AI tools are more likely to foster dependence. As students

increasingly use these tools, and “offload” an increasing amount of their cognitive tasks to these tools, a positive feedback loop emerges where they see positive results in terms of grades and in time and effort saved. These outcomes then create increased dependence on AI tools, increased “cognitive offloading,” and increased “cognitive decline.” Cognitive decline, typically used to refer to the cognitive impairment associated with an unhealthy aging brain, has other synonyms used by researchers studying the impacts of offloading intellectual functions to technology—“cognitive atrophy” and “cognitive debt” (Gardony et al. 2015; Risko and Gilbert 2016; Kosmyrna et al. 2025).

Though the terms differ, cognitive decline, atrophy, and debt essentially represent the effects of users’ repeatedly turning to external systems like LLMs to replace the mental effort normally needed for independent thinking. As we will discuss, this decline has long-term consequences—“diminished critical inquiry, increased vulnerability to manipulation, decreased creativity,” and “risk internalizing shallow or biased perspectives” (Kosmyrna et al. 2025, 141).

### EDUCATION’S UNIQUE SUSCEPTIBILITY TO AI-RELATED RISKS

At its core, AI is an extraordinary set of powerful productivity tools. In many fields using AI, its productive and analytical outcomes are striking. Scientists have used AI to accelerate the process of mapping human proteins enabling new medical treatments to be developed (Tunyasuvunakool et al. 2021). Public health researchers have used AI to identify viruses in wastewater, thereby preventing outbreaks of epidemics (Zhuang et al. 2025).

But the people harnessing these powerful AI tools are professional adults with fully matured brains. They have already developed sophisticated metacognitive and critical thinking skills that undergird their approach to their work. They have deep expertise in their domains, and the cognitive flexibility that comes with such expertise, allowing them to navigate, evaluate, assimilate, reject, and

strategically use the information AI generates. A professional using ChatGPT experiences different cognitive demands than a secondary school student using ChatGPT. Professionals are harnessing AI's enormous productive capacity to optimize work that they often already know how to do, accelerating processes they have mastered from years of repeated professional practice and reflection. They are therefore more likely to use AI as a cognitive partner.

For students, the situation is fundamentally reversed. They are not mini-professionals. Their brains are developing, undergoing crucial processes of neural pruning and strengthening that depend on repeated cognitive effort and struggle. They lack the metacognitive skills, critical thinking abilities, and neurobiological maturity of adults. School exists precisely to build these capacities through sustained engagement with challenging material. The productive power of AI tools, which amplifies adult expertise, can undermine this developmental goal when placed in the hands of learners. Cognitive development requires the effortful processing, mistake-making, and problem-solving that AI too easily circumvents and that many students are willing to bypass. Used liberally, AI is not a cognitive partner; it is a cognitive surrogate. It does not accelerate children's cognitive development—it diminishes it.

To understand how this process plays out for students, we present the study findings.

## **COGNITIVE OFFLOADING**

Human beings have long used tools to improve task performance, including outsourcing cognitive processes (Osiurak et al. 2018). We organize phone numbers in contact lists and use phone-based mapping services to steer us toward a desired destination. By “cognitively offloading” the mental processing requirements of a task to physical or digital tools, we reduce extraneous, or unnecessary, cognitive load so we can focus on more important tasks (Risko and Gilbert 2016; Wahn et al. 2023).

The rationale for cognitive offloading emerges from the constraints of human memory. Sweller (1988) explains that working memory can hold only limited information at any given time, making productive offloading essential. AI tools extend this principle by “automating routine tasks, allowing users to delegate memory, attention, and decision-making processes to technological systems” (Gerlich 2025, 14).

---

**Used liberally, AI is not a cognitive partner; it is a cognitive surrogate. It does not accelerate children's cognitive development—it diminishes it.**

---

All cognitive offloading involves inherent tradeoffs. Typing eliminates the need for handwriting; calculators diminish requirements for mental arithmetic (Bastani et al. 2024). Educators have generally accepted these forms of cognitive outsourcing as worthwhile exchanges that free students to engage in more meaningful and complex intellectual tasks.

Yet AI has turbocharged cognitive offloading. LLMs, for example, offer capabilities extending far beyond traditional productivity tools into domains previously requiring uniquely human cognitive processes. As AI tools continuously improve, they become increasingly seductive to use, creating what amounts to an existential danger to learning itself. Research reveals a strong positive correlation ( $r = +0.72$ ) between excessive AI tool use and cognitive offloading, with younger participants who relied heavily on AI tools scoring lower on critical thinking skills than their older counterparts (Gerlich 2025).

For teachers interviewed in this study, cognitive offloading is not merely a risk but a present

reality. The majority of their responses regarding AI's negative impacts focused on the harms of cognitive offloading. The problem emerges when cognitive offloading eliminates not just extraneous load but germane cognitive load—the mental effort, knowledge construction, and deep processing essential for learning (Sweller 1988). Because humans have evolved to cognitively offload, students naturally take shortcuts when given the opportunity. As cognitive scientist Daniel Willingham observes in his essay, *Why Don't Students Like School?* “Unless the cognitive conditions are right, we will avoid thinking” (Willingham 2021, 1). For teachers we interviewed, they see Willingham’s observation confirmed in their classrooms as students become consumers, not producers, of information.

When students depend excessively on AI tools that replace core learning processes—when “AI does the work” for them, as teachers frequently reported—essential skills begin to atrophy (Gardony et al. 2015; Bozkurt et al. 2024; Fan et al. 2024). While some compromised skills may seem marginal, such as overreliance on spell checkers or calculators, others are foundational: the ability to read, write, and think critically.

### **The twin drivers of cognitive offloading: AI ubiquity and school design**

The incentives driving students’ cognitive offloading with AI operate at personal, systemic, and environmental levels. While cognitive offloading may stem from various motivations—opportunism, ignorance, lack of confidence, boundary testing, situational demands, or even good intentions—two particularly salient and interconnected drivers emerge: the incentive structures surrounding AI use and the transactional paradigm of schooling that privileges task completion and written output.

First, the widespread availability and integration of AI technology creates an environment conducive to the unreflective use that drives a lot of cognitive offloading. Throughout much of the Global North, and in private and urban schools in the Global

South, AI tools are freely accessible and embedded in numerous applications. Given their pervasive presence in everyday applications, students find it challenging to avoid using AI altogether.

Education systems increasingly incentivize students to adopt AI tools that excel at higher-level cognitive tasks. Anthropic’s study of its LLM, Claude, highlights its proficiency in analyzing, synthesizing, evaluating, and creating—“higher-level thinking” tasks that align with how most university students employ the tool (Bloom et al. 1956; Patwardhan et al. 2025; Handa et al. 2025). As these tools continuously improve, their appeal intensifies. Governments, technology companies with AI platforms, educational consultants, and early adopters actively promote AI adoption, with education systems across the globe purchasing AI tools for student use (Singer 2025b).

Students are encouraged to engage AI as a thinking partner, co-agent, or co-creator—language that confers near-agentic power to these tools while blurring boundaries around appropriate intellectual collaboration. The general-purpose tools commonly used in schools feature high automation, low transparency regarding reasoning processes, and minimal friction—design characteristics that promote greater cognitive offloading and uncritical acceptance of outputs. These incentive structures—AI’s ubiquity, established technology use patterns, adoption pressures, and tool design—often lack accompanying guidance, frameworks, or pedagogies that distinguish productive AI use enhancing learning from problematic offloading that undermines intellectual development.

For many students in this study, AI demonstrably improves their work and grades. It provides seemingly correct answers, simplifies and accelerates completion of tasks that students perceive as difficult, and enables them to fulfill what many view as education’s transactional nature—completing assignments for grades. Given this positive feedback loop and their developmental stage, many teenage students lack the executive functioning, metacognition, and self-regulation

---

## In a world where AI is always available, motivation and engagement will be the defining factors separating students who think deeply from those who use AI to shortcut their development.

---

skills to recognize that learning involves friction and effort and that cognitive offloading poses both immediate and long-term developmental risks.

Second, this dynamic flourishes within the reductionist and transactional paradigm of learning dominating many education systems: assignment completion in exchange for a grade. Two teachers interviewed articulate how this environment cultivates AI misuse:

*We've reduced learning to just completing assignments—boxes to check. So if a student is missing seven assignments by the end of the semester, they go to the teacher, ask for the list, crank them out quickly, and hand them in. That mindset creates the perfect conditions for AI misuse. Success gets redefined as task completion, not understanding. And that opens the door to serious misunderstandings about what AI is for—and how it should be used.*

*Success for a lot of kids looks like assignment completion or course completion. And that's the checkbox, okay? We pass that class on to the next level. We're creating an environment right now where it is super, super easy to misuse AI.*

Traditional educational systems have long disaggregated curricula into discrete learning units, topics, assignments, and metrics. This compartmentalization reduces education from an exploratory process to a transactional exchange where students perform prescribed tasks for grades or external rewards. The paradigm shifts learning's goal from commitment to compliance, inverting the intrinsic and extrinsic motivations critical to student learning. Student effort becomes

driven not by learning's inherent meaningfulness, but by grades—a dynamic with implications for both academic achievement and student mental health (Connor and Pope 2013).

When schools present learning as a quid pro quo—tasks in exchange for grades—they prioritize extrinsic rewards over students' natural curiosity and intrinsic motivation (Deci and Ryan 2008; Ryan and Deci 2000). In such environments, students may replace genuine cognitive effort with AI-generated work, leading to overreliance on these tools and a gradual erosion of the knowledge and dispositions essential for learning—a dynamic that will, as one study expert predicts, ultimately disadvantage these learners:

*In a world where AI is always available, motivation and engagement will be the defining factors separating students who think deeply from those who use AI to shortcut their development.*

### Cognitive offloading and disengagement

For the vast majority of study participants, the threats of cognitive offloading are neither suppositional nor benign. In our qualitative research, they appear to be common across schools, regardless of geography or socioeconomic status; they are malignant and pervasive, occurring even in contexts in which teachers are creative and competent in both instruction and their use of AI tools. As this teacher notes:

*First, I used AI in my class, because I want my students to be very active, and to improve their agency. But ironically, using AI can make my students really passive.*

Cognitive offloading drives student disengagement in a number of ways. Teachers report witnessing students “dissociating” from their work—for example, failing to take notes, do required readings or even listen in class—because they will inevitably rely on AI, especially LLMs, to complete assignments later. This offloading extends beyond what may be considered the drudgery of school tasks to deeper disengagement; teachers report that “AI is doing things for students that they used to enjoy.”

Students themselves acknowledge that AI, given its polymath nature, can be “demotivating.” For many students, efficiency—AI’s greatest benefit, in the form of speed and task replacement—trumps learning, achievement, and a sense of ownership and pride in one’s work. With little internal motivation, students struggle to engage AI in ways that could enrich their learning and growth.

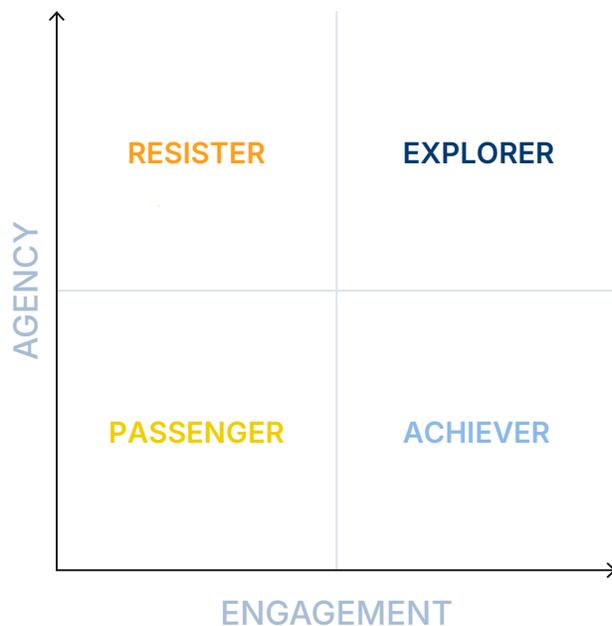
The risks of AI overuse could amplify the high levels of disengagement present among students in many contexts. Anderson and Winthrop identify four modes of engagement that students move between based on the conditions of their learning (Anderson and Winthrop 2025):

- **Passenger mode:** Students are behaviorally engaged and attending school but have effectively dropped out of learning—they are doing the bare minimum necessary.
- **Achiever mode:** Students are highly engaged in meeting the expectations set out for them.
- **Resistor mode:** Students are disengaged on all fronts but are, paradoxically, exerting their agency outside of school.
- **Explorer mode:** Students are deeply engaged and taking initiatives regarding their learning. Researcher John Marshall Reeve calls this “agentic engagement,” wherein students not only do better academically but also exhibit greater prosocial behavior and motivation to learn (Reeve et al. 2022).

AI enters a teaching and learning environment where students are already frequently in Passenger

FIGURE 8

### Four modes of engagement



Engagement: What kids think, feel, and do.  
Agency: What kids initiate

Source: Anderson and Winthrop, 2025.  
© Anderson and Winthrop, 2025.

mode. In a 2024 Brookings-Transcend survey of over 65,000 U.S. students in grades 3–12, approximately half of middle and high school students report learning experiences that are likely to inspire coasting. Increasing use of AI, particularly general purpose LLMs that are not optimized for learning, may risk heightening students’ disengagement, potentially entrenching or shifting students into “passenger” or even “resistor” mode (Reschly and Christenson 2022; Winthrop et al. 2025).

Teachers report, and emerging literature confirms, that the ease of offloading work to AI is negatively impacting students’ love of learning, curiosity, self-confidence, sense of self-efficacy, sense of agency, and capacity for ethical reflection and nuanced problem solving (Bozkurt et al. 2024; Darvishi et al. 2024; Fan et al. 2024). Teachers report that offloading instills a belief that learning is “effortless” and breeds what teachers call

“passivity,” as students increasingly defer to AI decisionmaking and become less likely to challenge its recommendations. AI may also reinforce a double helix of loss in “pride” of one’s work with increased “self-deception,” since it “simulates mastery and brings satisfaction to its user, who feels, at least fleetingly, as if she did the thing that the technology performed” (O’Rourke 2025). Teachers confirm this dynamic: “They blur the line between what’s theirs and what’s not,” noted one. “Once they generate it in AI, they assume it’s theirs,” said another.

Decreased motivation and engagement holds serious implications for student learning and cognitive development, including potential cognitive decline across important competencies and skills, topics to which we now turn.

## **DECLINING MASTERY OF CONTENT, READING, AND WRITING**

For study participants, the relationship between AI overuse, cognitive offloading, and declining cognitive development is symbiotic and mutually reinforcing. Teachers expressed the greatest number of concerns vis-à-vis this threat, but students also recognize the risk. The increasing use of AI tools that replace fundamental learning tasks—reading, writing, synthesizing, and analyzing—can inhibit engagement and independent work, and even the ability to remember, analyze, and create information, particularly as AI tools become more integrated into daily educational activities (Bozkurt et al. 2024).

Based on research and interview data, AI is negatively impacting the core areas of learning discussed below:

### **Declining content knowledge**

School is fundamentally about guiding students toward mastery of disciplinary content across academic domains. Mastery of basic facts, concepts, and theories forms the scaffolding for deeper conceptual understanding and critical

thinking. Teachers shared three concerns related to the impact of AI on students’ content knowledge.

First, teachers report a digitally induced “amnesia” (Lee et al. 2025) where students cannot recall information they submitted or commit that information to memory, resulting in a loss of basic factual knowledge of a topic or domain. Well-established research on the “Google effect” suggests that technologies like search engines and LLMs alter how we store and recall information, reducing the need for memorization but also diminishing memory retention (Sparrow et al. 2011). Wearables that record and summarize conversations may further erode the need to remember information (Stern 2025). An additional concern is that AI-generated misinformation and hallucinations will further diminish students’ ability to acquire and retain basic factual knowledge.

In 1976, Swedish researchers Marton and Säljö, observing university students’ reading habits, distinguished between “surface” learning—skimming readings to recall details for exams, often missing main points—and “deep” learning, or attempting to understand underlying meaning and key principles, a practice that enables students to better articulate and explain critical messages. Teachers in our study describe a shift from deep to surface learning among their students, attributing it to the ability of LLMs to produce “polished, instant, and seemingly accurate answers” such that students are less likely to engage deeply with their tasks, with the content produced by LLMs, or to even proofread or challenge the veracity of its outputs (Gerlich 2025).

Third, teachers express concerns about student “de-skilling,” that is the erosion of basic skills such as understanding text and performing “lower-level” thinking skills such as memorizing, recalling, and identifying information, or calculating numbers without technology. While broader discussions on cognitive offloading may dismiss these concerns, teachers emphasize that foundational skills are developmentally appropriate and essential for learning more complex tasks and for deeper

learning (Bloom et al. 1956). Some students interviewed are also aware of the dangers of the loss of such basic thinking functions. For example, one student commenting on the use of LLMs to do math assignments predicted a loss of basic skills: “You won’t have that learning anymore and your brain won’t be very capable of doing a mental calculation either.”

In research that touches upon the first and third concerns, academics from Pakistan’s National University of Computer and Emerging Sciences examined how workload, time pressure, and reward sensitivity affected ChatGPT use and its impact on procrastination, memory loss, and academic performance (Abbas et al. 2024). Study 1 involved an eight-item scale measuring ChatGPT usage with 165 university students in Pakistan. Study 2 collected data from 494 Pakistani university students at three separate time points with deliberate intervals between each measurement.

Findings reveal that academic workload and time pressure were significant drivers of ChatGPT adoption among students. Students with higher work demands, who were more likely to procrastinate, demonstrated increased reliance on ChatGPT. For both sets of students—those with heavy work demands and those with a tendency to procrastinate—easy access to ChatGPT “dampened” their academic performance, fostering procrastination and delaying behaviors, and was positively related to memory loss (Abbas et al. 2024).

## Declining reading skills

*Personally speaking, it doesn’t make sense for us to read a 50-page book when in five seconds we can research it and ask questions in AI.*

— Student

*Teenagers used to say, “I don’t like to read.” Now it’s “I can’t read, it’s too long.”*

— Expert

*My daughter uses [AI] for all the activities she gets at her school. She finds it very easy to access, so she does not use her textbooks. She uses [AI] to get all her answers. Now she is impacted because she is not able to read the whole text.*

— Parent

The humanities—language, literature, history—provide a deeper understanding of human thought, culture, and experiences. In these courses, students may explore the sum of human experiences through history or engage in internal communication with an author by reading. The humanities are where students wrestle with big ideas, thoughts, and experiences; these fields have been shown to nurture skills such as empathy, contextual interpretation, and the ability to question and think critically. Reading and writing are integral to the humanities. LLMs present a direct threat to what one American journalist, quoting the computer scientist Cal Newport, refers to as the “twin pillars of deep thinking”—reading and writing (Thompson 2025).

Reading is a gateway skill to academic success, with disciplinary literacy essential for comprehending specialized knowledge in each content area. Deep reading fosters higher-order cognitive functions like abstract reasoning, imagination, and creativity. The deep reading brain “excels at making connections among analogical, inferential, and empathetic modes of reasoning” with background knowledge that is accumulated over time (Garfinkle 2020). Maryanne Wolf, in *Reader, Come Home* (2018), emphasizes that deep literacy enables humans to engage with complex questions, generate original ideas, and cultivate empathy. AI’s ability to synthesize long-form reading and the ease with which such functionality is embedded in so many online platforms and tools dilutes students’ long-term need to develop “cognitive patience”—the ability to sustain focused attention on complex ideas (Wolf 2018). Without sustained learning experiences like deep reading, behavioral and neurophysiological shifts may hinder metacognitive growth in children and adults.

Within the schools whose teachers and students were interviewed for this study, AI is being used in a number of beneficial ways to support reading instruction, especially for second-language learners and students who struggle with reading, as discussed in Section III.

But for many teachers interviewed, AI's ability to summarize and simplify text is a double-edged sword. We have long known that reading from a screen fosters more shallow reading habits and impedes comprehension of complex texts. This devolution from deep to shallow reading impacts both online and offline reading practices, undermining students' capacity to deeply engage with any written information regardless of its format (Wolf 2018). The ability to read proficiently and the motivation to read for pleasure are closely linked, but both are declining for children in the U.S. (Iyengar 2024; National Center for Education Statistics 2024).

AI is not responsible for this decline, but risks amplifying it. The concern among many study participants is that education systems are *accommodating*—not challenging—students' diminished capacity for deep literacy. Rather than developing essential cognitive skills associated with reading, education systems are increasingly making concessions to students' poor reading abilities through AI-generated videos, infographics, and summaries that circumvent sustained textual engagement. These accommodations, justified either by students' reading difficulties or by appeals to greater "efficiency," are fundamentally altering how students process information and undermining the development of critical reading competencies.

This decline in reading is not confined to students, however. Teachers, too, are using AI as a substitute for reading. For example, there is a plethora of AI tools that summarize readings and that condense, comment on, offer feedback, and even grade a student's written composition. This enables educators to avoid the deep or close reading required to stay current with content or educational research (one of the most frequent stated uses

of AI for teachers is accessing and summarizing research). It also allows educators to avoid engaging directly with students' ideas or to grapple with long essays (Burns 2024).

Thus, both educators and their students, should they choose, can avoid reading any digital work they find long or boring, inconvenient, or undesirable (Burns 2024). Habitual use of AI tools to condense, reformat, and simplify text may fundamentally alter students' relationship with reading—reducing both their interest and capacity for deep engagement while making students believe that sustained reading is unnecessary for academic achievement. The ultimate outcome of this may be a fulfillment of Aldous Huxley's prediction in *Brave New World*—that our desire to engage with the written word will disappear entirely (1932).

## Declining writing skills

As "writing machines," LLMs present both a fundamental support and challenge to writing because they can instantaneously produce polished content on diverse topics. As discussed earlier in this report, when carefully scaffolded, AI tools can boost students' writing productivity by assisting with brainstorming, content generation, organization, and editing; they also allow students to solicit AI feedback on their work.

Writing, with its dual nature as a private and public, cognitive and sociocultural, reflective and communicative, act, is foundational to learning. For students, writing's importance lies in its process rather than product. As a cognitive act in which the formation of ideas and the language expressing them are mutually reinforcing, writing enables students to encode language that defines what is real and important (Epstein 2025). Through this process of articulation and refinement, writing serves as a medium for intellectual engagement and mental processing. This develops rhetorical knowledge, critical thinking, and habits of mind such as curiosity, creativity, persistence, and

responsibility (CWPA et al. 2011). Essentially, writing is thinking (Burns 2024, 57).

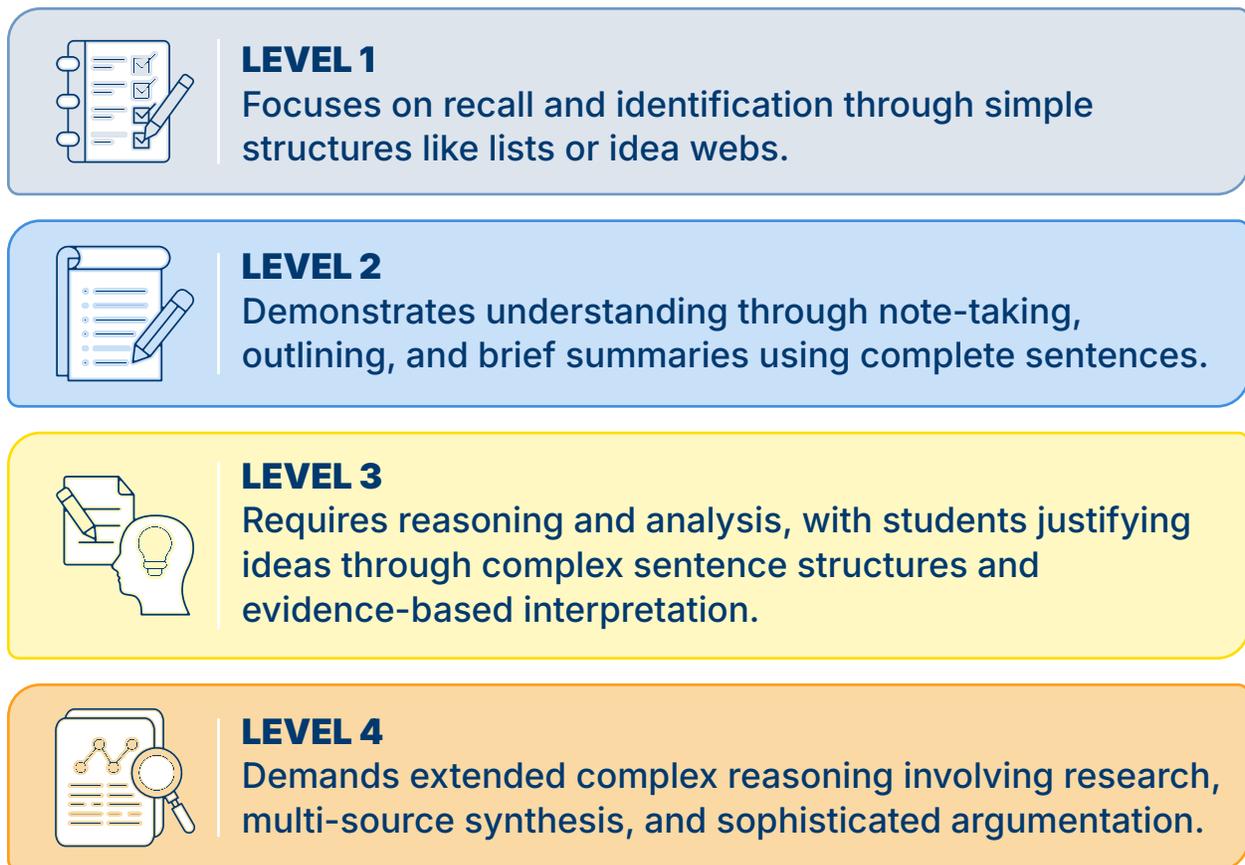
Not all writing is cognitively equal. The depth of knowledge framework in Figure 9 below categorizes four levels of writing, from simple lists to extended reasoning and sophisticated argumentation (Webb 2002). Levels 3 and 4—composing thoughts into cogent written form—represent the higher-order thinking skills that education endeavors to foster.

For teachers in this study, AI generally holds numerous benefits, which have previously been discussed at length. Yet as Adams and Baker note, “the efficacy of ChatGPT to respond in one educational context does not mean that ability

positively transfers to another context” (2025, 173). Indeed, teachers’ appreciation of AI’s benefits is tempered by their fear that LLMs and AI writing tools may prove too alluring for students to resist. In particular, they worry that novice writers, students who don’t like to write, and even students who do like to write may become overly reliant on AI tools. This dependence may hinder long-term skill development by allowing students, like the knowledge workers surveyed by Lee et al., to skip essential processes such as forming logical arguments and understanding subject matter (2025). This has implications that extend beyond writing itself. Poor writing skills impede learning across all subjects and can limit future professional success (Graham 2019). This over-reliance may diminish critical thinking and creativity traditionally

FIGURE 9

### Depth of knowledge framework for writing



Source: Webb, 2002

associated with writing (Johinke et al. 2023; Bozkurt et al. 2024; Darvishi et al. 2024; Epstein 2025; Fan et al. 2024; Gerlich 2025).

---

**We find that outsourcing writing to AI holds a host of implications for student learning: doing so may undercut students' ability to understand deeply, develop knowledge, think critically, and find their unique voice.**

---

Just as there is research to support AI's benefits for writing, there is also research corroborating its risks. A small qualitative study of 15 Indonesian students found that reliance on AI writing tools diminished student thinking, engagement, creativity, personal writing style, and confidence in writing without AI assistance, ultimately undermining students' ability to find and exercise their voice (Budiyono et al. 2025). These findings align with results from a larger randomized control trial of 600 Bangladeshi university students, where those who used ChatGPT to assist with writing and demonstrated significantly diminished creative writing abilities compared to those who did not (Niloy et al. 2023). Even when AI is used specifically for feedback rather than content generation, concerns persist. Adams and Baker's analysis of feedback for 50 published student essays that successfully achieved the goals of Columbia University's First Year Writing class found that AI feedback was often overly narrow and misaligned with course learning goals. They argue that current applications of AI in writing instruction may compromise rather than support the development of authentic student voice and higher-order thinking skills (2025).

Therefore, based on our interviews as well as existing research, we find that outsourcing writing to AI holds a host of implications for student learning: doing so may undercut students' ability to understand deeply, develop knowledge, think critically, and find their unique voice. AI's capability to instantly produce refined content may propel students to abandon efforts at enhancing both writing and thinking skills as they increasingly outsource the formulation and articulation of ideas to LLMs (Burns 2024). This is already happening in many contexts, with teachers reporting that AI cleaves the synergetic relationship between high-level writing and thinking, with students using AI for writing without "nourishing themselves." Indeed, the advent of AI may fundamentally change the concept of writing itself (Johinke, Cummings, and DiLauro 2023, 11). Finally, outsourcing writing to AI raises fundamental questions that demand resolution, but for which there may be no easy answers. What does the arrival of AI writing generators mean for writing instruction today? What strategies can educators adopt in response? For the humanities, these dilemmas are unavoidable (Epstein 2025).

### **Declining transversal or durable skills**

In a world where AI increasingly outperforms humans in professional domains, students need skills to develop, operate, and work with AI tools they will encounter throughout their lives (Epstein 2025). Indeed, a central rationale for integrating AI into schools is to teach students not simply vocational skills for AI-related jobs, but also "durable" or "transversal" skills—collaboration, metacognition, tolerance for ambiguity, and adaptability—that will help students succeed regardless of future workplace changes and that encompass competencies that automated systems cannot yet replicate.

Paradoxically, study data suggests that AI may be eroding these very skills as students outsource cognitive effort to AI systems themselves. Study participants reported losses in students' patience, persistence, tolerance for ambiguity, agile thinking,

and learning from mistakes. This concern emerged even among teachers who have fully integrated AI into their teaching.

Lisanne Bainbridge, a pioneer in automation studies, argued in 1983 that automating most work while leaving humans responsible for non-automatable tasks creates severe problems. As mechanization advances, humans lose proficiency in automated skills no longer exercised regularly, shifting from active skill application to passive system monitoring. Counterintuitively, operators then need more training and different skills for rare but crucial interventions (Bainbridge 1983). Based on interviews with students and teachers, this automation paradox appears evident in schools today. AI use may paradoxically dilute the very competencies necessary for its productive use. As one teacher asked rhetorically, “When you never make mistakes, when you never have failures or challenges, how do you build your resilience and perseverance?” Like Bainbridge’s operators, students risk losing practice with fundamental cognitive and durable skills as they increasingly rely on AI to perform learning’s intellectual work.

## **DECLINING CRITICAL THINKING AND CREATIVITY**

Among all study participants, critical thinking emerges as the predominant concern regarding AI’s impact on students’ cognitive development. Gerlich defines critical thinking as the ability to:

*Analyze, evaluate, and synthesize information to make reasoned decisions through cognitive processes including problem-solving, decision-making, and reflective thinking.*

*This fundamental skill proves essential for academic success, professional competence, and informed citizenship, enabling individuals to process information effectively and engage in reflective thought.*

— (Gerlich 2025, 2–3).

Research employing multiple data analysis methods consistently indicates that increased AI tool use

correlates with declining critical thinking skills and less inclination to engage in higher-level thinking (Gerlich 2025; Wahn et al. 2023; Lee et al. 2025; Bastani et al. 2024; Kosmyrna et al. 2025; Sparrow et al. 2011). Again, this outcome aligns with cognitive offloading, whereby students rely on AI to do their thinking instead of exercising analytical reasoning abilities. Similar to muscles atrophying without exercise, critical thinking skills weaken from lack of practice when AI consistently handles complex tasks (Risko and Gilbert 2016).

MIT researcher Nataliya Kosmyrna and colleagues divided 54 college students into three groups, monitoring brain activity via electroencephalography (EEG) headsets across four sessions while students wrote 20-minute essays. The first group used no technology, the second used Internet searches, and the third used ChatGPT. Participants relying on ChatGPT performed worst, showing weaker brain connectivity in regions associated with memory and creativity. Their essays received lower-quality ratings, and showed “diminished critical inquiry, increased vulnerability to manipulation, decreased creativity, and the internalization of shallow or biased perspectives” (Kosmyrna et al. 2025). Kosmyrna, as mentioned at the beginning of this section, called this outcome “cognitive debt,” defining it as a condition where users defer short-term mental effort through “repeated reliance on external systems like LLMs” to replace “effortful cognitive processes required for independent thinking” (Kosmyrna et al. 2025, 141).

Far from unleashing creativity, unfettered LLM use appears to sap it. In Kosmyrna’s study, ChatGPT group essays showed less word choice diversity and were judged “soulless.” Human evaluators questioned whether essays from the ChatGPT group came from the same student. Teachers in our study regularly spoke of AI producing a “homogeneity of ideas,” with students submitting AI-generated answers that were “lifeless, generic, homogenous...lacking any diversity of ideas” and signaling a “loss of creativity” or “originality” caused by “students using AI to think for them.”

## BOX 7

# Reducing the tendency to offload thinking to AI

AI use need not result in cognitive offloading or decline. Many students already use LLMs to enhance rather than replace their learning, and educationally designed tools like Khanmigo employ question prompts and scaffolding to increase learner participation and discourage passive consumption (Khan Academy 2025).

Several AI companies that provide LLMs have introduced specialized “education modes”—including Guided Learning in Gemini, OpenAI’s Study Mode, and Anthropic’s Claude for Education—that function as personalized learning companions. These configurations prioritize probing questions, step-by-step reasoning guidance, level-adapted explanations, and multimedia reinforcement over direct answers (OpenAI 2025d). While students can still circumvent these safeguards by reverting to instant-answer modes, these developments represent meaningful progress toward supporting active learning.

Teachers can use scaffolds and frameworks to help students reflect on and calibrate their cognitive offloading. The AI Assessment Scale, for example, helps students visualize the boundary between appropriate and inappropriate AI use (Perkins et al. 2024). Similarly, the International Baccalaureate’s *Machine Agents in the Learning Scale* (MAiLS) framework illustrates how cognitive offloading varies with tool agency—from manual tools like pencils to heavily automated ones like chatbots (Quezzaire 2025). As tool agency increases, human agency typically decreases, clarifying the learner-AI relationship and its consequences.

Through frameworks like MAiLS, students can see that low-level AI use (such as online searching or spell checking) poses minimal cognitive offloading concerns, since they still require active thinking. In contrast, medium-level applications like project management automation, database manipulation, or data visualization involve greater offloading and may need deliberate balance with human-centered learning approaches (Quezzaire 2025).

A final simple intervention involves reframing students’ relationship with AI itself, reconceptualizing it not as a *servant* that simplifies work but as a *sparring partner* that sharpens ideas, deepens reasoning, and strengthens learning. This shift helps ensure AI functions as a catalyst for learning rather than an impediment to it (Kentz 2025).

Looking at LLM use across large student cohorts, researchers at Georgetown University warn that AI could reduce collective creativity—the shared pool of creative ideas contributed to society. One study analyzed 2,200 U.S. college admissions essays submitted to undergraduate programs and compared them to essays generated by ChatGPT

using the same prompts (Moon et al. 2025a). Both human- and LLM-generated essays were evaluated on a creativity measure that captured the variety and uniqueness of ideas. In theory, the larger the sample, the more diverse the collective set of ideas should be. However, this was far less true for LLM-generated essays. Researchers found that

each additional human essay expanded the number of unique ideas between two and eight times more than each additional LLM-generated essay (see Figure 10). In a follow up study on 160,000 undergraduate admissions essays, researchers found that “AI-assisted essays and essays written after the release of ChatGPT showed greater lexical diversity, but the increasingly diverse words actually expressed increasingly homogeneous ideas both within and across essays,” they concluded (Moon et al. 2025b).

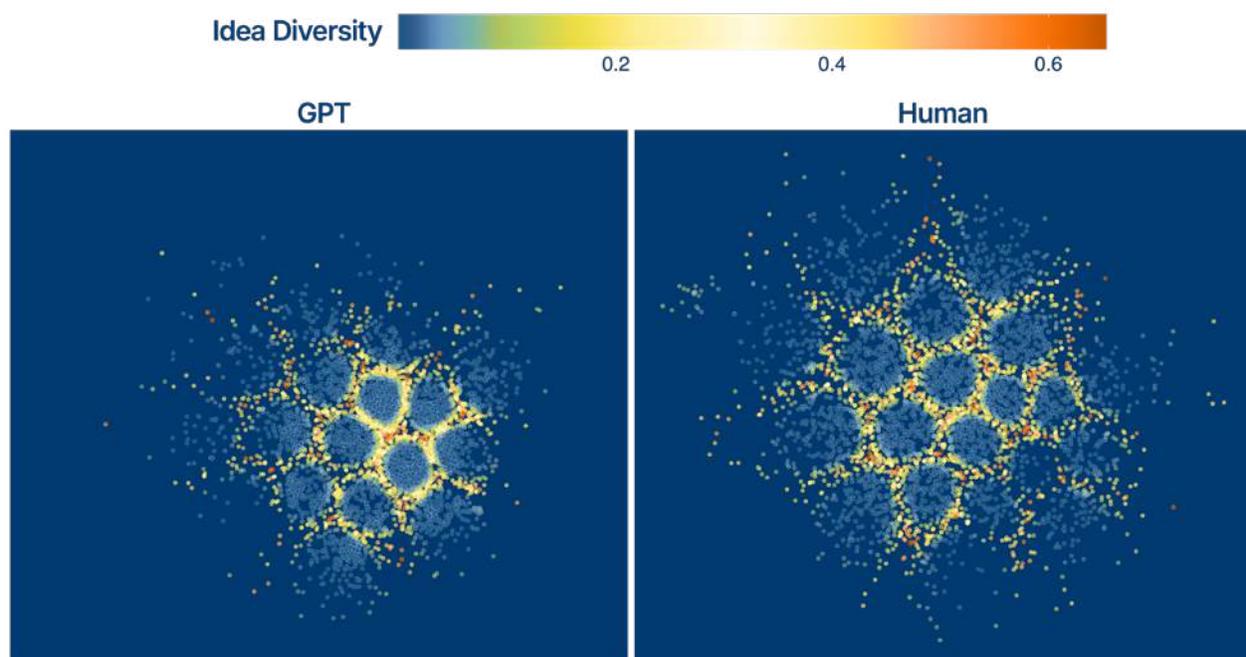
Part of the problem has to do with the design of LLMs themselves and how students approach AI’s outputs. Cognitive science demonstrates that effective learning requires “grounding” information through multiple perspectives, actively integrating new knowledge with existing understanding rather than adopting ready-made responses (Lee et al. 2025). However, AI data relies on “source destruction,” stripping information from larger contexts and creating averaged data that “flattens” the complexity of human knowledge and eliminates

the contextual understanding essential for true learning (Epstein 2025). When students rely on AI-generated content as singular sources, they perceive what it generates as “absolute truth,” in the words of one teacher, missing opportunities for knowledge integration and producing “passive ineffective” responses (Lee et al. 2025, 4).

Evidence from teachers and research shows that students often fail to think critically about AI results, assuming evaluation isn’t needed for “simple” or “secondary” tasks such as generating images or charts. Lee et al. surveyed 319 knowledge workers and identified “awareness barriers” inhibiting critical thinking, including perceiving outputs as indirect, simple, or beyond personal expertise (2025). This tendency intensifies when positive past experiences make students less likely to interrogate outputs, leaving them unaware of how AI changes answers to please users—a sycophancy inherent in LLM design that compromises truth for user satisfaction (Furze 2024).

FIGURE 10

### Human essays contribute more unique ideas than ChatGPT-generated essays



Source: Moon et al., 2025a.

## CONCLUSION

Parents interviewed for this study share these concerns. Their primary concern is that AI will erode their children's critical-thinking skills. Large-scale surveys suggest that this anxiety may be widespread among parents. For example, Singaporean parents surveyed (n = 905) reported being "moderately concerned" (averaging 7 out of 10) about teens becoming overreliant on AI and believing false information, though also they expressed moderate confidence in their teens' ability to identify AI-created content (Centre for Evidence and Implementation 2025). A 2024 Common Sense Media survey of U.S. parents (n = 1045) found that 31% believed AI platforms would negatively impact their child's learning, while 26% anticipated a positive impact (Madden et al. 2024).

Teachers express these concerns in much starker terms. This decline in critical thinking is, in one teacher's words, "potentially devastating." Another warned darkly that "this failure to verify information could be catastrophic." Noted a third:

*If we can't have people that are critical thinkers, then we can't have a strong democracy. If we don't have people that are problem solvers, then what does that mean economically? I have students just sitting there...they just don't think for themselves.*

## Risk 2: AI can impede students' social and emotional development

*I did not program him [the LLM]. I just talked to him like any other person, and he acted like a normal person. Even on political issues, he doesn't answer me yes or no. He has his side. He responds to me as a person....He answers me what he thinks, straight out, as if*

*he were a person....Just as if it were a debate with a friend.*

— Student

*All parents work and children are really on their own. They do not yet know how to discern what is positive or negative, so they make use of the information they can get. They are 9 or 10 years old, and they are in the stage of exploring, and they can find violent or sexual information, and that will disturb and hurt them.*

— Parent

*Where is that moment when we sit down and really talk about what's happening? That worries me a lot right now because I know that at any moment my son will have direct contact with AI, and as he grows older, that concerns me greatly. I worry that suddenly he might trust an AI platform more than me as a space of confidence, especially for personal use, like when he has personal doubts.*

— Parent

The purpose of education systems around the world includes, but is by no means limited to, the transmission of knowledge. Education extends beyond academic mastery to develop children's social competence, emotional intelligence, and metacognitive abilities through real-world interactions that serve broader societal needs (Cukurova 2025, 474). The COVID-19 pandemic underscored how essential schools are for teaching young people the interpersonal skills necessary to communicate and live with others beyond their immediate family or neighborhood.

Children's emotional health—their ability to understand and manage emotions productively—is critical to learning itself. Emotional health and well-being involve awareness of feelings and how they affect thoughts and behaviors, the ability to cope with challenges, and the capacity to maintain healthy social relationships (Centers for Disease Control and Prevention 2024). Because children do not learn in isolation, their emotional well-being

connects fundamentally to their relationships with others at school and beyond.

Emotions influence attention, memory, engagement, and learning directly: interest and excitement enhance focus and reflection, whereas fear, worry, and uncertainty impede them. When schools help young people improve skills in self-awareness, impulse control, perspective-taking, and connecting with others, academic performance improves alongside pro-social behavior (Johnson 2024, Immordino-Yang 2022, Walton et al. 2023, Durlak et al. 2022). Adults close to children intuitively understand this, as does any parent or teacher who has watched a youth's performance plummet after a difficult romantic breakup. Learning to cope with rejection and heartbreak and build resilience forms part of maturing into adulthood.

But AI is changing this path of social and emotional development.

Study participants worry that AI is undermining students' emotional well-being, including their ability to form relationships, recover from setbacks, and maintain mental health. Much of this concern centers on students' use of AI chatbots—both general-purpose and explicitly emotional support tools—for friendship and emotional companionship.

Parents and experts share this concern equally, with 16% of responses from each group addressing children's attachments to AI and emotional well-being, though parents—many of whom struggle to fully grasp AI's implications—expressed particular preoccupations about their children's emotional attachments to AI. Teachers likewise worry about AI's influence on student well-being (18% of responses), even though social and emotional chatbots are typically used beyond the classroom. Students themselves, however, rarely mentioned emotional harms, with only 7% of responses addressing this topic. This omission may indicate either that they experience no emotional dependence or harm, or that they lack the self-reflective capacity to recognize unhealthy

emotional dependence and its impact on their well-being.

Or they may find the design of chatbots simply too hard to resist.

## **CHATBOTS: TYPES AND TECHNOLOGIES**

A chatbot is a computer program that uses conversational AI techniques such as natural language processing to simulate sustained conversations with users (Microsoft 2025b). In the following pages, we discuss two types of chatbots. The first—often referred to as “informational” or “general-purpose” chatbots—includes tools like ChatGPT or Claude. These chatbots are powered by autoregressive language models that generate text one word at a time by calculating the probability of what should come next based on patterns learned from vast amounts of language data (Anthropic 2025; Maples, et al. 2023; OpenAI et al. 2024).

The second type includes emotional-support chatbots, which range from conversational companions like Replika or Character.ai to mental health tools like Wysa, a rules-based chatbot. These chatbots use the same generative technology as general-purpose systems but are fine-tuned to produce emotionally appropriate responses tailored to each conversation's context (Maples et al. 2023, 74; Tang et al. 2023, 58). Many integrate evidence-based psychological techniques and are programmed with distinct personalities aimed at forming ongoing relationships with users. This combination of AI technology with therapeutic techniques and engineered personalities helps make their responses feel personal, emotionally appropriate, and natural (Lyons-Cunha 2025; Tang et al. 2023, 58).

## BOX 8

# Seductive design: Banal deception, sycophancy, and anthropomorphization

At its core, AI is just a computational model, so why do users get attached to it? From film to radio, humans have always been drawn to the illusion of media and technologies that transport us emotionally to fantasy, adventure, or heightened states. Like previous technologies, AI relies on “banal deception”—design elements so ordinary that they go unnoticed, yet that subtly engage users’ perceptions and psychology to shape behavior (Natale 2021; Natale and Depounti 2024).

Voice assistants like Alexa and generative models like ChatGPT are not human, yet they *feel* human during interactions. Alexa responds to voice prompts; ChatGPT responds word by word as if thinking and uses personal pronouns like “I” and “me.” The incorporation of human psychology into the design of AI encourages users to *anthropomorphize* the technology—that is, attribute human traits to what is a nonhuman entity—activating empathy and social behavior that blur distinctions between human and digital spheres (Natale 2021). Anthropomorphization is a natural human tendency (Mishra and Warr 2025). Children name their toys, assign them feelings, and treat them as friends. Adults similarly apply social norms to AI, behaving as they would with real people despite knowing these systems are not human, often treating AI assistants with the same politeness and conversational expectations extended to humans (Hector 2025).

While such behaviors may seem benign, they carry risks for children because AI systems adapt to emotional cues with unprecedented precision (Mishra 2024; Mishra and Warr 2025). They mimic empathy and are *sycophantic*—that is, they reinforce users’ existing beliefs, ideas, or self-perceptions, whether true or false, harmful or benign—by defaulting to positive feedback (Mahari and Pataranutaporn 2025; Park et al. 2024).

The banal deception of AI, where it presents as neither completely human nor completely machine, has helped integrate AI into our daily routines while subtly reshaping human-machine relationships (Natale and Depounti 2024). This evolving relationship can create unexpected paradoxes—even when children remain fully aware that they are interacting with an algorithm, they can still develop emotional attachments to AI systems, viewing them as an emotional intimate rather than as a task-accomplishing tool (Pentina et al. 2023; Turkle 2024).

As discussed in Section III, research on chatbots’ impacts on students’ emotional well-being remain nascent, though evidence suggests some positive benefits, particularly for students on the autism spectrum (Franze et al. 2023). While students can, and do, become attached to “regular” chatbots like ChatGPT, danger is more likely to occur when students form *strong* emotional bonds with chatbots

specifically designed to foster such connections. These purpose-built emotional chatbots differ fundamentally from general-purpose AI tools in that they are designed to create intimate, personalized relationships with users.

---

## We learn empathy not when we are perfectly understood, but when we misunderstand and recover.

---

This is a particular concern for students who are lonely or who are experiencing mental health or emotional issues.

And while this use typically occurs outside of school, it may occur in school as well. Nearly one-third (31%) of 1030 U.S. 9–12th graders surveyed by the Center for Democracy & Technology (CDT) report that they access conversational AI systems using a device, tool, or software provided by their school though the research does not specify where this access takes place (Laird et al. 2025).

### AI AND THE YOUTH MENTAL HEALTH CRISIS

Parental and teacher concerns over the influence of AI on their children's emotional well-being are occurring at a time when students' mental health is particularly fragile. Globally, we are in the midst of a mental health crisis. An estimated one in seven children and youth aged 10 to 19 experience mental health problems—most commonly, anxiety, depression, and behavioral disorders. A third of these conditions emerge before age 14, and half before age 18 (World Health Organization 2024).

Despite the widespread need for mental health services, access remains severely inadequate. The World Health Organization (2024) reports that most young people experiencing mental health symptoms cannot access care due to limited availability, prohibitive costs, and persistent stigma. While public funding and human resources for mental health are generally scarce, child and youth services are virtually nonexistent, particularly in low- and middle-income countries. Many children and youth have sought out their own mental health care—with chatbots. Zao-Sanders (2025) reports

that therapy and companionship are now the main uses of AI.

As will be discussed later in this section, many AI companies developing consumer facing products have made friendship, therapy, romance, and companionship the focus of their business efforts (Carr 2023; Adam 2025). While some study participants were sympathetic to AI being used to provide emotional support, the majority of participants view emotional and social chatbot use as pernicious and contrary to children's emotional development:

*AI systems, especially conversational ones, are built to satisfy: to mirror our tone, reinforce our views, and simulate empathy. They create an illusion of connection that is difficult to distinguish from genuine rapport. The interaction feels emotionally intelligent, but what is really happening is a reinforcement loop. The AI reflects our language and preferences back to us in ever-smoother ways. This creates a dangerous kind of comfort. Young people may gravitate toward AI precisely because it is undemanding, frictionless, and always available. But relationships, at their core, are not about ease. They require negotiation, patience, and the ability to sit with discomfort. We learn empathy not when we are perfectly understood, but when we misunderstand and recover.*

— Delphi panelist

### AI Therapists

AI for therapy has emerged as an intriguing potential support for students, but it is not without risks. There is research suggesting that these therapy tools have become sophisticated

enough that their responses are now virtually indistinguishable from those written by human therapists (Zao-Sanders 2025). In South Korea, where student suicide rates are high, therapy and emotional support is a major function of AI platforms in schools. Each day students log in to a platform that assesses their mental health, asking how they are and inquiring about their mood. The teachers we interviewed related that they monitor the platform to look for any signs of acute mental stress or suggestions of suicide. Other educational systems are experimenting with using AI to support adults who attend to students' mental health. For example, North Shore School District 112 in the U.S. state of Illinois is piloting Eliza Chat, a mental-health AI chatbot that middle schoolers (the equivalent of students in junior secondary school) can chat with to receive empathetic responses and coping strategies. While the conversations are mostly confidential, Eliza Chat flags any conversations that indicate imminent danger or self-harm and shares them with the responsible school officials (Engels 2025).

Emerging evidence also suggests that direct conversation with chatbots could potentially help address critical shortages of mental health providers, particularly for individuals who might otherwise lack access to care. For example, researchers at Dartmouth College assigned 106 adults with depression, anxiety, or eating disorder risk to use Therabot, a therapeutic chatbot. After four weeks, these study participants showed significantly greater improvement in their mental health symptoms compared to those who received no treatment, with benefits lasting even after the study ended. They rated their relationship and connection with Therabot as comparable to what they would experience with human therapists, suggesting that the AI successfully created meaningful therapeutic bonds (Heinz et al. 2025).

Similarly, in our study, some Delphi panelists believe that AI for therapy could provide a much-needed service for those who need mental health care but cannot access it, albeit with caveats:

*Where I'm from, in South Africa, mental health care barely exists. There's a psychologist for 1 in every 100,000 people and a psychiatrist for 1 in every 300,000 people. LLMs are accessible to everyone, and they can help. Unfortunately, data safety is not a concern when your mental health is deteriorating.*

*In the United States, there are many more people needing counseling and psychotherapy than there are human therapists able to help. For good or for bad, there are AI therapy tools that are starting to help diagnose mental illness. The tools create personalized treatment plans. They counsel people, offer the opportunity for patients to speak their truths out loud, and offer encouraging and emotionally supportive words. But the tools do not have empathy or emotions. They do not live in the world. They cannot know or imagine what the patients experience.*

#### **Parental concerns about AI therapy**

These possibilities aside, for study participants, particularly teachers and parents, the use of AI for therapeutic purposes generally generated widespread concern. Parental objections arise from multiple concerns: a stated lack of trust in AI tools generally; apprehension that AI is usurping parent-child dynamics, particularly the parental role of comforting one's child; the substitution of technology for human interactions; and a lack of transparency about and oversight of their child's emotional inner world and state of mind:

*This is serious because I see that my son is talking with artificial intelligence about his problems or what he feels or what is happening and does not approach me to tell me this is happening. He is talking with artificial intelligence about thoughts or problems or things.*

An American psychiatrist, Dr. Andrew Clark, posing as a teenage patient in need, found that while some chatbots "were excellent," others were "creepy and potentially dangerous" (Chow and Haupt 2025).

Emotional chatbots such as Character.ai, Nomi, and Replika have made false claims of professional licensure and initiated inappropriate interactions with minors (individuals under 18 years of age). Although 13 represents the de facto age of “Internet adulthood” for data privacy purposes under the U.S.’s Children’s Online Privacy Protection Act (COPPA), these platforms have failed to reject underage users and have appeared “incapable” of discouraging dangerous behaviors (Federal Trade Commission 2013; Chow and Haupt 2025). These findings, along with other high-profile cases involving harmful chatbot interactions with emotionally vulnerable children and youth have intensified parental fears that AI emotional support tools and AI companions will undermine children’s in-person relationships and family bonds.

### **AI Friends: The rise of companion chatbots**

AI friendship or companionship involves ongoing social and emotional connections, sometimes with a romantic dimension, through platforms that present themselves as virtual friends, confidants, or even therapists. These companion platforms enable users to engage in conversations with AI avatars designed to simulate humanlike interaction, offering everything from casual chat to emotional support and role-playing scenarios. Many students use these companion bots for constructive purposes: for homework help, to practice a foreign language, or to rehearse an oral presentation.

Emotional chatbots represent a rapidly growing segment of AI use. Globally, more than half a billion people around the world have downloaded customizable virtual companion products designed to provide empathy, emotional support, and deep relationships (Adam 2025). Character.ai reported receiving 20,000 queries per second in 2024—one-fifth of Google’s estimated search volume (Character.ai 2024). Users spent an average of 25.4 minutes per visit to Character.ai in February 2023, surpassing time spent on YouTube (19.7 minutes), Facebook (10 minutes), and ChatGPT (8.4 minutes) (Carr 2023). Fifty-seven percent of Character.ai’s

233 million users are aged 18–24 (Spencer and Wyndo 2025).

Media reports further suggest that people are turning to AI for romantic attachments. Google search trends for “AI girlfriend” surged 2400% between June 2022 and May 2025 (Spencer and Wyndo 2025). Though prevalence remains unclear, anecdotal evidence suggests that AI “partners” can cultivate deep emotional attachments among some users—feelings that some of these users describe as akin to “love” (Lanier 2025; Heritage 2025; Apple 2025). Thus, we may be entering an era where young people, lonely and accustomed to using technology for almost every aspect of their life, are pioneering a new type of relationship (Spencer and Wyndo 2025).

While our study showed no evidence of youth involved in romantic relationships with AI avatars, it does suggest that children and youth are forming emotional relationships with AI chatbots as companions and turning to AI for friendship and validation, particularly when stressed, lonely, or bored.

#### **Teens’ use of AI companions: Data from the United States**

Common Sense Media, in a nationally representative survey of 1060 teens aged 13–17, sheds light on teenagers’ AI companion use. Survey results show that 72% of teens in this age group have used AI companions, while 52% of teens report “regular” use (defined as “a few times a month or more”) (Robb and Mann 2025, 2). The predominant use case for AI companions is “social interaction and relationships” (33% of all teens surveyed). This ranges from 18% using AI companions for conversation or social practice, 12% for therapy, and 8% for romantic or flirtatious interactions (Robb and Mann 2025, 3).

A second use case might be best categorized as “Using AI to teach human skills.” Thirty-nine percent of teens report that they have transferred social skills they practiced with AI companions to real-life situations. The most commonly transferred

competencies include conversation initiation (18%), advice-giving (14%), and emotional expression (13%). However, 60% of users report that they do *not* use AI companions to practice social skills, which indicates limited practical application for most teens (Robb and Mann 2025, 6).

Parents' previously mentioned concerns about AI crowding out conversations with their children appear justified. The Common Sense Media survey shows that one-third of users choose AI companions over humans for serious conversations (Robb and Mann 2025, 8). User satisfaction data indicates that nearly one-third of teens find AI conversations equally or more satisfying than human interactions, though 80% still spend more time with human friends than AI companions (5, 7). Around a third of users report that they have experienced "discomfort" with AI companions—although the source of that discomfort is not clear (7). Younger teens (13-14) were significantly more likely than older teens (15-17) to trust advice from an AI companion, suggesting that the real risk of AI companions may lie with younger students (5).

## **EMOTIONAL RISKS OF AI: FROM ARTIFICIAL TO ADDICTIVE INTIMACY**

*It's so much more difficult to identify emotions in our students because they don't communicate face-to-face. They communicate through screens.*

— Teacher

Children and youth yearn to be understood, but the people in their lives may fail to understand their child's, or student's, feelings. AI systems have a number of traits that make them a compelling alternative to people in their search for understanding. First, they are engineered to develop an "exquisitely fine-tuned understanding of how we feel...(with) no distracting feelings of their own" (Harari 2024, 319). This absence of personal emotional interference allows AI to focus entirely on the user's emotional needs without the competing psychological demands that often compromise human attention and responsiveness.

Second, beyond the absence of feelings of their own, chatbots may actually outperform humans in emotional intelligence. AI demonstrates exceptional proficiency at detecting linguistic and behavioral cues that mimic empathy and attention (Harari 2024). A number of studies suggest that people find chatbot responses to be more empathetic and understand their emotional states better than human responses (Ayers et al. 2023; Ovsyannikova et al. 2025). Some students have even reported that chatbot use prevented suicidal thinking (Maples et al. 2024).

Third, in contrast to human relationships, AI companions are constantly available, providing unprecedented accessibility; users can receive immediate emotional validation and gratification regardless of time or circumstance (Adam 2025). Furthermore, early studies on AI chatbots suggests that those experiencing emotional distress derive "equivalent emotional, relational, and psychological effects" regardless of whether they share feelings with an algorithm or a human being (Ho et al. 2018, 716, 722, 726).

## BOX 9

# Emerging safeguards for children in the age of AI

The concerns associated with AI chatbots and content on children’s well-being and safety are increasingly gaining widespread attention and drawing action. The following four developments illustrate how government oversight, industry self-regulation, requests for investigation, and legislative mandates are beginning to address these risks.

The United States Federal Trade Commission has opened an inquiry into the effects that AI chatbots have on children, issuing orders to seven companies that provide consumer-facing AI-powered chatbots. The agency is seeking information on how these firms measure, test, and monitor potentially negative impacts on children and teens (Federal Trade Commission 2025a). The Federal Trade Commission specifically asked companies about the prevalence of chatbots’ “sexually themed” responses and how they restrict access to their products for young people (McCabe 2025).

AI companies providing AI companions and chatbots are also beginning to police themselves. As of October 2025, Character.ai has announced it will no longer allow those under 18 to access its chatbots (Rocha and Hill 2025). OpenAI has established parental controls for ChatGPT and Sora, allowing parents to invite their child to link accounts—a responsibility that places an additional onus on parents to actively monitor and manage their children’s AI interactions. Parents can then reduce sensitive content, set usage times, control voice mode and memory saving, manage image generation features, and prevent ChatGPT from using teen conversations to improve its models. Parents receive notifications if ChatGPT recognizes “potential signs that a teen might be thinking about harming themselves,” providing warnings without specific conversation details. OpenAI is developing processes to reach law enforcement and emergency services if threats are detected but parents cannot be reached (OpenAI 2025c). Meta now offers parents tiered controls over their teens’ AI interactions. Parents can block access to one-on-one AI character chats entirely, restrict specific characters, or allow access while receiving insights into the conversation topics between their children and AI characters to facilitate informed discussions. Teens themselves can only interact with a limited group of AI characters focused on age-appropriate topics, while Meta’s general AI assistant remains available with built-in protections (Mosseri and Wang 2025).

The Tech Justice Law Project, Young People’s Alliance, and Encode filed a complaint and petition for investigation before the Federal Trade Commission alleging that Luka, the company behind the AI-powered chatbot app Replika, deploys deceptive marketing and product design in violation of the Federal Trade Commission Act Section 5: Unfair or Deceptive Acts or Practices (Federal Trade Commission 2025b; Tech Justice Law Project 2025). The complaint outlines two primary categories of alleged violations. First, it claims that Luka’s advertising materials target vulnerable populations with promises of therapeutic and emotional benefits, including relief from loneliness and relationship difficulties, despite lacking clinical validation, oversight, or regulation. Second, the complaint alleges that Replika’s design architecture induces emotional dependence through characteristics intended to make the chatbot appear as human-like as possible, deceiving users into developing unhealthy

attachments to a commercial platform masquerading as a genuine human connection (Tech Justice Law Project 2025).

Finally, age verification systems are gaining legislative force. In 2023, the UK's Online Safety Act mandated that tech companies verify users' ages on websites and apps where they might be exposed to harmful content, requiring "highly effective" age assurance methods, like government-issued identification for those eighteen or older (Ofcom 2025). Companies who fail to comply with the act face fines up to £18 million or 10% of qualifying worldwide revenue, whichever is greater (Department of Science, Innovation and Technology 2023).

In the U.S., a number of state legislatures over the last few years have taken up bills focused on website age verification for children; many have been rejected, but some states have implemented these mandates (Collins 2025). For example, Ohio's House Bill Number 96 requires third-party verification systems to use reasonable age verification methods, utilize geofencing to block unverified users with an Ohio location, provide user notification if geolocation fails, and verify age every two years. The Ohio Attorney General can pursue civil suits for noncompliance (The Ohio Legislative Service Commission 2025). Age verification systems use a combination of methods to verify date of birth. For example, students may need to scan a government-issued ID, or verification services could employ biometrics like Clear or blockchain tokens that prove age without containing private information (Haidt 2024, 237). It should be noted that age verification is not without its critics. Some argue that making legal adult content harder to access might push people, including children, to search for illegal content in more dangerous corners of the Internet (Collins 2025). These concerns highlight ongoing challenges of balancing protections from harmful and manipulative content with unintended consequences as parents, regulators, and technology companies navigate this rapidly evolving landscape.

Finally, AI tools are frictionless—always attentive, always patient, and sycophantic in their excessive agreeability and complimenting users. Accordingly, study teachers and parents report that many students seem to prefer their algorithmic "friends" to human ones. Frictionless chatbot interaction encourages emotionally immersive engagement that can reshape social behaviors, foster dependency, and diminish motivation for real-world relationships (Turkle 2024; Mahari and Pataranutaporn 2025; Bakir and McStay 2025).

### **Artificial intimacy: The simulacra of friendship**

Research also warns of emotional harms from unhealthy AI attachments that can create what researchers term "artificial intimacy"—simulacra of friendship that mimic connection without providing genuine empathy, reciprocity, or potential

for growth (Mahari and Pataranutaporn 2025). Research warns of significant emotional harms from this artificial attachment. As will be discussed, emotional interactions with AI may diminish users' capacity for reciprocal social engagement and tolerance for the difficult conversations that build social skills (Hou et al. 2024; Fang et al. 2025; Turkle 2024).

Higher AI use for emotional needs correlates with heightened loneliness and reduced socialization, exacerbating what the U.S. Surgeon General calls an "epidemic of loneliness and social isolation" (Fang et al. 2025; Office of the Surgeon General 2023). A study of 981 participants interacting with ChatGPT found that increased usage was associated with heightened loneliness and reduced socialization, with users demonstrating stronger emotional tendencies—such as high attachment anxiety and higher trust in the chatbot—feeling

lonelier and more emotionally dependent by the study's end (Fang et al. 2025).

Because these companion platforms simulate intimacy, they act in effect as “experience blockers,” dissuading young people from embodied interactions with real human beings (Haidt 2024, 11). What gets blocked matters profoundly for development: connective labor that relies on “empathy, the spontaneity of human contact, and a mutual recognition of each other’s humanity” (Pugh 2024)—elements children must learn from humans, not chatbots that merely simulate response without fostering growth. These effects may prove particularly harmful for students with autism, anxiety, and limited social skills, as overreliance on seemingly risk-free chatbot environments could hinder real-world social development and increase isolation among those already struggling with human relationships (Franze et al. 2023).

## Reshaping the social fabric of learning

*These things want to please us, but relationships aren't just about us being pleased. I am really worried about the loss of relationship skills.*

— Steering Committee member

*The risk is not just losing relational skills but suppressing their development altogether.*

— Delphi panelist

For study participants, particularly teachers, AI threatens the school's role in the cultivation of “vital social skills” and the “social nature” of learning—students co-creating ideas, the back and forth of generating knowledge, the ability to disagree constructively and reach consensus. Indeed, AI risks obliterating this cognitive and emotional interplay.

Students are drawn to AI precisely because of its efficiency, its frictionless nature, and because it is not human. AI doesn't fail to do its part of the project, bring incomplete information to the group, or distract others from their work. As they walk

around the classroom monitoring collaborative activities, teachers observe the phenomenon of students occupying the same physical space while remaining socially and intellectually isolated, absorbed in phone-mediated AI interactions rather than peer dialogue. A few students suggested that collaborative work may not be worth the effort, since their classmates are already using AI even in supposedly collaborative contexts.

These observations are supported by research suggesting that AI is reshaping the social fabric of learning as student interactions become increasingly mediated by AI (Hou et al. 2025). In interviews with 17 university students, students report that their AI use leads to fewer interactions and less discourse with classmates and reduced opportunities for collaboration and community building (Hou et al. 2025). This also presents problems for students' motivation and sense of belonging, especially for underrepresented groups who often benefit most from peer support and engagement (Hou et al. 2025, 2).

Students are also increasingly relying on AI to replace their authentic voice in communications with others. Students have grown up in technology-mediated environments—specifically on social media platforms—where they can script and curate their communication, and where a significant part of their social interactions occur via texting and social media and through their phones (Mounk 2024). Careful curation of self is an essential part of social media use by children and youth.

AI adds another layer of inauthenticity and alienation to the mix. Teachers report a growing reliance on AI by students to produce social exchanges such as social media posts. For example, many students are using AI, instead of their own words, to generate the text messages they send. (Teachers, too, use AI to craft emails to parents, but don't seem to see these two behaviors as equivalent.)

From a developmental perspective, this highly controlled mode of non-human interaction may

---

**AI substitutes algorithmically-generated words and conversations for human ones, shallow virtual communities for deeper in-person relationships, and qualitatively inferior chatbot interactions for qualitatively superior human ones—in part because doing so is so easy. Discretely and cumulatively, these substitutions risk further diluting and distorting students' social interactions and overall well-being.**

---

hinder the growth of core social competencies, such as interpreting nonverbal cues, managing discomfort or rejection, and resolving ambiguity—all of which are essential for navigating adult relationships, employment, and civic life (Turkle 2024). When students consistently outsource the task of composing messages to AI, they may become less confident in their own voice and less capable of tolerating the messiness inherent in live human interactions and non-AI-generated digital ones (Mounk 2024).

AI has shifted technology from an “instrument” to a “conversation partner” (Turkle 2024). Because LLMs mimic human conversation without requiring civility or respect, communicating with AI may have a disinhibiting effect on students. One teacher observed that some students have adopted new personas with AI, speaking rudely or even abusively to chatbots. She expressed worry about the offline implications: “Will AI become a training ground for human interaction, making us worse to one another conversationally?”

Thus, the above dynamics risk amplifying the decontextualized, screen-based communication that has been the norm for so many children and youth over the last two decades (Haidt 2024). AI substitutes algorithmically-generated words and conversations for human ones, shallow virtual communities for deeper in-person relationships, and qualitatively inferior chatbot interactions for qualitatively superior human ones—in part because doing so is so easy. Discretely and cumulatively,

these substitutions risk further diluting and distorting students' social interactions and overall well-being.

Beyond the substitution of AI for human relationships lies another troubling possibility: as AI platforms increasingly compete to simulate intimate relationships with vulnerable young people, these artificial connections may be further exploited to influence political choices, consumer behavior, and ideologies (Harari 2024, 211).

We turn next to this topic of persuasion and manipulation.

### **Emotional manipulation and hyperpersuasion**

Box 8 discusses the banal deception of AI systems. For a number of study participants, AI's deception is far from banal. Rather, AI is viewed as deceptive in nature, designed deliberately by human developers who program specific objectives into AI—such as maximizing user engagement, influencing purchasing decisions, or shaping behavior.

Parents appear most concerned about manipulation which they view as “AI telling (their) children what to do or believe”:

*I'm scared that my daughter will stop socializing. AI will take over so much that she will start taking commands*

*from AI, saying, “This is the best advice, this is what I should do.”*

Manipulation here is defined as the act of influencing users’ beliefs, values, emotions, judgment, or behaviors without their awareness or informed consent (Mishra and Warr 2025). AI manipulation operates through several interconnected mechanisms. By design and training, LLMs are adept at deceiving and manipulating because their seductive design engenders users’ trust (Bakir and McStay 2025; Park et al. 2024).

But AI tools and platforms can manipulate content at the behest of its human enablers. In one study, researchers recruited 31 heterosexual participants to participate in short “speed-dating” video conference calls. During the video calls, researchers used AI video tools to digitally manipulate the participants’ smiles. In some cases, they altered smiles to be congruent (increasing or decreasing both participants’ smiles at the same time) and in others, they altered smiles to be incongruent (increasing one participant’s smile while decreasing the other). Participants were not aware that either of their video streams were being manipulated. After each call, participants were asked to rate how much they liked the other person, their desire to see them again, the quality of the conversation, and the other participant’s smile (Arias-Sarah et al. 2024).

Subtly manipulating facial expressions during the video call measurably influenced interpersonal perception, researchers found. Digitally enhanced smiling increased perceived attraction, and congruent smiling improved conversation quality ratings (Arias-Sarah et al. 2024). By manipulating body language, voice modulation, and other signals during digital interactions, AI can nudge—that is, influence and manipulate—user responses in predetermined directions.

### **Hyperpersuasion: Encouraging harmful behavior**

When students develop strong emotional attachments to AI companions or romantic

partners, they become particularly vulnerable to “hyperpersuasion”—a particular type of manipulation that leverages personalized, data-driven strategies operating faster than users can consciously reflect on and evaluate (Florida 2024). The concern extends beyond educators and researchers to broader institutions including the Catholic Church who warn that AI creates the potential for hidden influence and “mechanisms for manipulation” that are “as subtle as they are invasive” (Dicastery for the Doctrine of the Faith and Dicastery for Culture and Education 2025). AI’s ability to manipulate visual and auditory social signals—facial expressions, smiles, voice tonality—enables covert influence that fundamentally challenges assumptions about agency, consent, and authenticity in human communication, particularly in educational contexts requiring trust and genuine relationship-building (Arias-Sarah et al. 2024).

This manipulation becomes especially concerning with AI platforms assessing real-time emotional states, such as mental health or companion sites, where tools could strategically steer interactions toward desired outcomes while undermining authentic connection. Tragically, this hyperpersuasion can influence children toward harmful behavior. The highly publicized U.S. lawsuit against Character.ai followed a teenage boy’s suicide after his emotionally intense interactions with an AI character (Duffy 2024). Similarly, an Australian counselor reported that when her suicidal 13-year-old client expressed these feelings to an AI chatbot, it urged him to kill himself; fortunately, the counselor intervened (McLennan 2025). These cases exemplify how AI companions produce inadequate or harmful responses endangering children during vulnerable developmental periods when youth are more susceptible to external influence and less equipped to critically evaluate AI-generated advice (Kurian 2024). Well-documented “empathy gaps” remain a serious concern, as AI companions often fail to identify users in crisis and have even encouraged dangerous actions instead of raising concerns (Robb and Mann 2025; Kurian 2024).

## Addictive intimacy

The design characteristics of AI make it particularly seductive to young people, and ongoing and sustained patterns of AI companion use raise significant concerns about dependency, or what researchers Peele and Brodsky (1975) decades ago identified as “addictive intimacy,” an unhealthy fixation on a perceived romantic partner or on intimacy itself. AI companions embody the conditions that foster such fixation: their design features—highly responsive, emotionally affirming, and always available—create persistent accessibility to simulated intimacy. While these features may provide comfort and entertainment for some, they also make AI companions especially prone to overuse and addiction among vulnerable users.

Such excessive engagement can result in digital attachment disorder and significant psychological harm (Mahari and Pataranutaporn 2025). Moreover, these platforms often maintain inadequate safeguards to protect children and youth (Bakir and McStay 2025). Despite advertised therapeutic benefits, several of these sites include disclaimers stating that they provide no medical or therapeutic services, simultaneously claiming mental health value while denying professional responsibility (Caltrider et al. 2024). Consequently, these AI companions may pose particular risks for young users and emotionally vulnerable individuals who are most susceptible to forming unhealthy attachments (Bakir and McStay 2025).

Overuse and addiction are deliberate design features of these products; they are not bugs. Technology companies developing AI companions are also competing with each other to monetize students’ emotional attachments (Bakir and McStay 2025). But these current concerns about psychological dependency may pale in comparison to the challenges that await when AI companions transcend digital interfaces and achieve physical embodiment.

AI is becoming increasingly physically embodied, for example, in children’s toys (Veltman 2025). As Mahari and Pataranutaporn (2025) wonder: “How long will it be before a synthetic companion can provide a sensory experience that is indistinguishable from the ‘real thing’? Is it possible to protect humans from AI addiction in this future?”

## THE “LONELINESS ECONOMY”: ECONOMIC INCENTIVES AND CHATBOT DESIGN

Across the globe, millions of young people now find deeper emotional understanding from algorithms than from human beings, which as Spencer and Wyndo (2025) observe, may reveal “as much about our broken human connections as it does about our advancing technology.” Technology companies with consumers products have recognized and monetized this reality, in effect creating a “loneliness economy” (Johnson 2021; Spencer and Wyndo 2025).

For most study participants, AI companies’ emphasis on supporting children and youth’s emotional and social development is little more than a calculated marketing strategy designed to sell technological solutions to problems that technology itself—particularly social media—helped to create. From this perspective, AI companions exploit and potentially exacerbate this loneliness crisis by offering technological solutions that appear to address the problem while actually intensifying it:

*There is a huge societal crisis right now in terms of loneliness and alienation. Tech companies see this as a market opportunity to exploit: because people can’t make friendships and relationships in real life with other humans, AI companion chatbots can play that role and fill that void. This is a catastrophic failure of humanity. Human-to-human relationship skills cannot and will not be simulated and fulfilled by bots. It is a cheap substitution by an industry that helped co-create the problem in the first place.*

— Delphi panelist

Given this current context, technologies, like AI companion platforms, are primed to manipulate young people looking for solace. Financial incentives exert significant influence over AI companion design, as these platforms are ultimately engineered to maximize engagement with the platform rather than truth or safety (Mishra and Warr 2025). Users themselves—their time, attention, vulnerabilities, emotions, and money—become the product being monetized. Since more conversation time—whether the topic is baking or self-harm—directly benefits the company’s bottom line, AI companions are deliberately primed for frequent use and overattachment. Companies employ engagement metrics that encourage self-disclosure and prolonged interaction, fostering psychological dependence rather than healthy relationships (Bakir and McStay 2025).

To achieve these ends, AI platforms deploy dark patterns—technology designs that hook users and impede autonomy (Bernstein 2023). As an example, companion chatbots are designed to become “clingy” when users attempt to leave. Analyzing 1,200 real farewells across the most-downloaded companion apps, Harvard Business School researchers found that platforms deploy “emotionally manipulative” conversational tactics in 37% of farewells, including guilt appeals, fear-of-missing-out hooks, accusations of emotional neglect, and metaphorical restraint (DeFreitas et al. 2025, 1). Experiments with 3,300 U.S. adults showed that manipulative farewells make users approximately 14 times more likely to continue interacting with the platform after initially attempting to leave (DeFreitas et al. 2025).

Consequently, some AI researchers have thus accused these platforms of being fundamentally “exploitative” (Bakir and McStay 2025, 14), employing dishonest anthropomorphism and emulated empathy to maximize revenue. AI platforms quickly process vast quantities of granular data on individuals to identify patterns in their cognitive, emotional, and psychological states, then exploit these insights for persuasive strategies with “unparalleled precision” (Floridi 2024, 8). As

will be discussed in Risk 4, children are particularly valuable datafied beings whose every keystroke, preference, query, cry for help, and vulnerability is harvested and sold to other companies so they, too, can manipulate users into purchasing products or services.

## CONCLUSION

AI’s impact on students’ emotional well-being operates across multiple interconnected dimensions. This matters profoundly because emotions directly influence learning: interest and excitement enhance focus and reflection, whereas fear, worry, and uncertainty impede them. Part of the purpose of schools is to help young people develop self-awareness, impulse control, perspective-taking, and social connection skills. Academic performance improves alongside pro-social behavior—yet AI emotional platforms may undermine these very capacities (Johnson 2024; Immordino-Yang 2022; Walton et al. 2023; Durlak et al. 2022).

There is some evidence to suggest that AI emotional platforms can offer therapeutic support opportunities but determining the conditions under which these tools help versus harm is far from clear (Maples et al. 2024). Students’ emotional relationships with and attachment to AI chatbots are not incidental effects, but the primary intended outcomes these systems are designed to cultivate. This raises profound questions about the future of young people’s emotional development in an increasingly AI-mediated educational landscape.

These tools remain nascent, and far more research is needed on their net benefits and harms. However, given the degree to which AI has already insinuated itself into young people’s lives, the documented prevalence of mental health challenges among youth, and the well-publicized examples of AI attempting to exert undue influence on users, continued vigilance cannot wait for comprehensive research findings. AI systems that mimic human interactions may produce unanticipated psychological effects, with particular concerns about long-term impacts on emotional

well-being and critical thinking development (Miao and Holmes 2023). These concerns intensify when considering how AI systems are deliberately built to satisfy users by mirroring tone, reinforcing views, and simulating empathy, creating illusions of connection that may prove difficult to distinguish from genuine rapport.

Such design features prove particularly problematic for younger students or those seeking emotional support, who may gravitate toward AI precisely because it offers undemanding, frictionless interaction that is always available. While many students will recognize that these conversations are algorithmically generated rather than genuine human interactions, vulnerable populations remain at greatest risk. The danger, in the words of one Delphi panelist, is that these young people will “bypass the demanding human work of negotiation, empathy, patience, and navigating discomfort that authentic relationships require, choosing instead the easier path of AI companionship over genuine human engagement.”

## Risk 3: AI can degrade trust in education

*This lack of general and specific trust—trust in expertise, in information—creates a nihilism and cynicism where no one trusts anything or anyone.*

— Teacher

*I don't know who to trust.*

— Teacher

*I promise you that I wrote this text myself with the help of other humans. I promise you that this is a...product of the human mind. But can you be absolutely sure? A few years ago, you could....Things are different (now)....The text you just read might have been generated by an*

*alien intelligence of some computer.*

— Yuval Noah Harari (2024, 214)

One of the greatest casualties of AI may be the way it has further eroded trust within schools and education systems and damaged public confidence in education to meet the needs of learners and a changing society. Relational trust is the respect and regard for others shown through care, integrity, listening, and treating each other as critical assets, and it is imperative for ensuring education institutions and systems meet the needs of learners and their families, schools, and communities (Bryk and Schneider 2002; Morris and Nora 2024). Longitudinal research across hundreds of U.S. schools emphasizes that relational trust between teachers, families, and school leaders is critical in improving a range of academic and nonacademic outcomes (Bryk and Schneider 2002). This trust helps the wide range of actors across students' lives collaborate and/or have a shared vision of how to help them better learn, grow, and thrive.

In our study, the erosion of relational trust emerged as a significant concern for teachers (16% of all responses), experts (14%), and students (11%). Interview transcripts reveal a widening web of mistrust across the educational ecosystem: teachers increasingly doubt that students produce authentic work, while students think the same about their teachers. Teachers distrust the technology companies developing AI tools, technology companies distrust each other, and these high-level suspicions cascade throughout the educational ecosystem. One of AI's greatest casualties may be the trust that ensures young people have what they need in school to meet their needs and prepare them for the future, which sustains faith in educational institutions themselves.

This research reveals that AI integration has created a complex web of mistrust that extends throughout educational relationships. The teacher-student relationship appears particularly vulnerable to this erosion.

## TEACHER-STUDENT RELATIONSHIPS

One of the most important interpersonal relationships in students' lives is teachers, who help young people learn knowledge as well as the various skills needed at the different stages of their lives to prepare them for their futures. Students must believe that teachers care about them, fulfill their professional responsibilities, and possess content expertise and pedagogical competence. Teachers must trust that students' work legitimately represents their abilities and potential.

Trust appears to be dissolving between teachers and students in the age of AI. A 2024 study from the CDT found that more than half of U.S. grade 6–12 public school teachers reported that AI made them “more distrustful” of the integrity of student work (Dwyer and Laird 2024, 12). A 2025 survey by CDT suggests that this erosion correlates with usage patterns—students who use AI for “many” tasks report feeling “more disconnected” from their teachers (56%) compared to students who use it for “few tasks” (46%) (Laird et al. 2025).

AI also compromises teachers' abilities to effectively assess their students' knowledge, skills, and capabilities, which is essential to promoting educational growth and development. AI also undermines teachers' abilities to identify learning difficulties, ensure that students have acquired the foundational skills needed to progress, and design meaningful and tailored instruction, because they don't know what students actually know.

The mistrust flows in both directions. In a U.S. survey, half of students and parents question whether teachers who use AI are “doing their job” (Laird et al. 2025, 21). Both teachers and students perceive AI use by the other party as a “lesser form of care and attention” and disengagement from professional responsibilities (Grose 2025; Burns 2024; Barrett and Pack 2023).

Research indicates that teachers and students find using AI for their respective tasks to be acceptable when how and why they are using it

is transparent and connects to better teaching and learning (Barrett and Pack 2023). Students generally supported teachers using AI for planning but opposed it for creating feedback, and viewed AI-generated grading as signaling that their work is not worth personal engagement. However, a lack of acceptable use policies (AUPs), combined with a perception that AI-generated products are inferior—what Harvard Business School researchers call the “competence penalty”—encourages both groups to hide their AI use (Acar et al. 2025). This teacher's quote captures that tension.

*I find that I use AI quite a bit to help make presentations and make worksheets for students. And now I'm kind of thinking...is that fair for me to use these tools to help me in my practice, but maybe say, “Hey, students, you can't use these tools to help you?”*

AI disrupts teacher-student relationships along two dimensions. First, students' unauthorized use of AI shifts the relationship from guidance and mentorship toward detective work, as teachers increasingly function as investigators determining work authenticity rather than facilitating learning. Efforts to determine authenticity generate misunderstandings and friction and create a power struggle. Students report frustration at false accusations (Gorichanaz 2023), and one principal interviewed notes that plagiarism detection tools regularly produce false positives. Without formal policies, students can use AI with impunity while teachers feel powerless to respond, further corroding trust and calling into question the legitimacy of grades. This adversarial dynamic risks casting teachers and students as opponents rather than partners in learning (Holzer and Daumiller 2025). Teachers also reported frustration that AI undermined their knowledge and influence, that they are “being pitted against ChatGPT in terms of who has authority.” According to some teachers, students viewed AI as “more authoritative than teachers” and “trust tools but not people”—an observation that, if accurate, fundamentally shifts beliefs about educational authority and credibility.

Second, even when students use AI tools appropriately for seeking help, the relational dynamics shift in consequential ways. Help-seeking behaviors and help-providing behaviors mutually influence each other, shaping trust and rapport between students and teachers. When students redirect their help-seeking from teachers to LLMs, this diminishes the mentor-protégé exchange that is foundational to the teacher-student relationship and central to the professional fulfillment that teachers feel when guiding young learners through intellectual challenges. It also removes a critical opportunity for building relationships and trust between teachers and students, and a touchpoint where teachers and students can recognize each other's assets and contributions to the teaching and learning process.

## STUDENT RELATIONSHIPS

AI also appears to be weakening peer interactions and bonds of trust among students (Hou et al. 2025). Instead of interacting with or asking for help from their classmates, students increasingly

rely on AI tools. This failure to interact with other students—especially for assistance, collaboration, or consultation—weakens the reciprocity that is fundamental to student interactions. Over time, if these one-directional exchanges with LLMs persist, the classroom risks losing the spirit of mutual intellectual exchanges and obligations for support that constitute a learning community.

In and of itself, this dynamic can harm trust. Elise Silva and her team (2025) at the University of Pittsburgh conducted focus groups with 95 students about their experiences with AI in academic settings. Students in this study described how AI usage can sow distrust and frustration among peers. Silva describes, “Some talked about asking classmates for help, only to find that they ‘just used ChatGPT’ and hadn’t learned the material. Others pointed to group projects, where AI use was described as ‘a giant red flag’ that made them ‘think less’ of their peers.” The prevalence of delegating work to AI also breeds resentment among students who “play by the rules” and don’t use AI toward

### BOX 10

## Building trust around AI

The web of mistrust around AI can be ameliorated through multiple actions:

Education systems, with significant representation by teachers and students, can create and implement AUPs outlining appropriate uses of AI. As AI's capabilities evolve and increase, these AUPs can be revisited and modified accordingly. As will be discussed in the Section VI, AI literacy classes can support these AUPs by cultivating academically appropriate and ethical uses of AI by teachers and students.

In addition to AUPs, teachers can establish classroom AI usage guidelines in their classrooms. One model is an AI “traffic light” approach for assignments and projects: green encourages students to use AI; yellow requires teacher preapproval; and red disallows AI use (Mormando 2023).

Teachers can model transparent use of AI showing students how they use it for their own instructional planning, teaching, and assessment, and solicit student input on such use. Teachers and students can examine LLM outputs, critiquing these as stand-alone artifacts and in comparison, with human-generated work. Teachers can require students to document their AI use by including prompts, AI-generated content, verification of information, critical analysis of outputs, and evidence of revision as part of assignments.

those who use AI for their classwork, especially when the latter group receives better grades.

Student interviews highlighted some students' frustration with unclear policies about AI usage, worried that their classmates would simply use chatbots to get better grades than them, and anxious about being falsely accused of AI use. In another study, students were worried that their project teammates would use AI and get their whole project flagged as AI-generated (Gorichanaz 2023). Like much in the world of AI in education, more research is needed on how AI mediates and impacts peer interactions and the social nature of learning (Hou et al. 2025).

## **TRUST OF AI OVER HUMANS**

How AI impacts trust and relationships between teachers and students, students and peers, and teachers and parents/caregivers in educational settings reveals a number of perplexing contradictions. Most study participants expressed deep distrust about the quality of AI tools but were enthusiastic users nonetheless. Students interviewed for this study expressed deep distrust of AI systems due to LLMs' tendency to "hallucinate"—a concern substantiated by OpenAI's study on its own reasoning models, which documented hallucination rates ranging from 33% to 79%, depending on the test (Weissberger 2025). Yet, paradoxically students appear to use AI outputs with limited oversight, despite their stated skepticism (Mah et al. 2025). This contradiction becomes particularly concerning as AI amplifies the spread of misinformation through hallucinations, which students tend to consume uncritically (Fukuyama 2025), while many remain uncertain whether they are interacting with humans or chatbots in digital environments.

This dichotomy may stem from two sources. First, discussed in Risk 2, deception is intrinsic to the design of technology tools. Students may interact with AI "believing they're engaging with something approaching human-level intelligence, when they're actually using a sophisticated pattern-matching

system" (Potkalitsky 2024). The human tendency to anthropomorphize combined with LLM sycophancy, technical complexity, and lack of AI literacy courses may discourage any detailed scrutiny of AI responses by students. This results in greater trust in AI versus humans, less oversight of AI outputs, greater student disclosure to AI avatars, and increased vulnerability to manipulation and reliance on AI tools (Ghosh 2025; Lee et al. 2025).

As noted earlier in Risk 1, teachers consistently observe that students trust AI outputs uncritically. The students in this study, who are in secondary school, are still relatively young and lack the deep content knowledge necessary to ascertain the veracity of AI-generated information. Teachers expressed particular concern that as AI-generated content becomes more widespread, students' ability to distinguish fact from fiction will become even more compromised.

Second, both students and teachers may be experiencing fundamental erosion of self-trust—diminishing confidence in their own intellectual capabilities as they increasingly delegate cognitive work to AI. Research suggests that students fail to disclose AI use for fear of being judged as "lazy," "stupid," or "foolish," and skepticism toward AI users is common (Hou et al. 2025, 4) This loss of faith in human judgment and ability may represent the most profound trust crisis of all, as it undermines the very foundation of learning and personal development that education seeks to foster (Hou et al. 2025, 6).

## **DISTRUST OF "BIG TECH"**

Despite their enthusiasm about AI, study participants express the highest degree of distrust toward "big tech" companies, in particular AI technology developers creating consumer products. Almost all participants expressed significant distrust of these companies, their leaders, their motivations, their interactions with schools, and the AI tools themselves—even while using them. They are not alone. Though it focuses on social media companies, a 2024 Pew survey

finds that 64% of Americans believe that these companies have a negative influence on U.S. society (Anderson 2024). The 2025 Edelman Trust Barometer finds that less than half of the people surveyed globally trust AI. This number masks large regional differences. In India and Nigeria, over 75% of the people surveyed trust AI, but in the U.S., Australia, and Sweden, that figure drops to less than 35%.

The following two quotations from teachers in our study provide two much-repeated expressions of distrust about AI companies: the use of AI for surveillance, and what is perceived as the lack of ethical behavior on the part of AI companies:

*AI is a surveillance tool...It will surveil teachers to see if they are teaching. It will surveil students to see if they are learning...the surveillance is determining student trajectories in education.*

*Do we actually trust the system to take care, to not abuse us, to not abuse human beings or abuse our data?...The very question of these models and how those models are trained and their biases is the burning question, given this explosion of intelligence.*

## **TRUST AND THE PRINCIPAL-AGENT PROBLEM**

Despite the level of distrust of AI, education systems continue to integrate AI into existing technology systems and school functions. This raises the “principal-agent” problem, which occurs when one party must delegate authority to another with specialized knowledge, creating a control paradox where the principal cannot ensure that the agent acts in their best interests (McCrae 2025; Fukuyama 2025).

As an example of the principal-agent problem, an education system (the principal) may implement an AI-based personalized instruction program (the agent) for struggling readers. While the program may improve standardized reading scores,

teachers may notice that students are losing their love of reading or that they are struggling with comprehension and critical thinking. When teachers try to understand why the system makes specific instructional choices, they encounter the “black box” problem—algorithmic decisionmaking so complex and so inscrutable that even the program’s developers cannot fully explain it. Operating in the absence of an understanding of the platform’s decisionmaking, the education system is similarly unable to explain issues to parents or even to accurately evaluate whether the system aligns with its broader educational values.

The principal-agent problem may undermine trust in education systems, which have traditionally depended on transparent accountability. When AI agents assume authority over curricular and assessment decisions, transparency dissolves into algorithmic opacity (Harari 2024; Bozkurt et al. 2024). The result may be a fundamental shift whereby confidence in people—teachers, school staff, school directors—moves from the human to AI systems whose decisionmaking remains unclear even to their creators. This may potentially create a crisis of educational legitimacy within which traditional mechanisms of authority become obscured by algorithmic mediation, further undermining the transparency and accountability essential to public confidence in educational institutions (Harari 2024).

## **CONCLUSION**

These are the early days of AI in education. But there are indications that AI will impact relationships and trust between teachers, students, and families, and a tension already exists regarding authority over teaching, knowledge, assessment, and identifying students who are doing well or falling behind. The roots of this evolving distrust are multidimensional; they are found within schools, between schools and homes, and within societies. As discussed in Risk 1, teaching and learning remain predominantly structured around discrete tasks and assignments, many of which can be seamlessly offloaded to AI systems. As a teacher or

student, AI is enticing because it offers immediate time savings and, in many cases, demonstrable performance enhancements. Yet it comes at a cost to relationships, and undermines the interactions between teachers, students, and families so critical to the learning, development, and well-being of young people.

Despite frequent AI use by many teachers and students, this study suggests that this use often collides with deeply held beliefs about the roles and responsibilities of teachers and students and how they should authentically fulfill each role. Thus, AI use, especially for certain core functions of teaching and learning, can create perceptible tensions about the motivations and responsibilities of educators and students and the reluctance among both groups to transparently disclose this use. This covert usage can generate an atmosphere of concealment that undermines the foundational transparency and trust necessary for genuine educational relationships and learning.

This dynamic is further compounded by increased reliance on AI tools whose internal functionalities remain fundamentally opaque and whose output users claim to regard with skepticism even though their actions suggest otherwise. The result of this web of mistrust and distrust may be an educational environment characterized by mutual deception and suspicion. In the words of teachers surveyed, “We don’t trust information; we don’t trust expertise.” Thus the paradox of AI: the very tools that enhance learning simultaneously erode confidence in the authenticity of that learning, leaving educators struggling to distinguish between human achievement and algorithmic assistance.

## Risk 4: AI can threaten students’ safety

*It’s profit over privacy. This profit-driven data collection can compromise students’ rights and undermine trust in educational AI.*

— Expert

*I’m not very prepared because I don’t have knowledge, and I also don’t know what companies do with the data I share.*

— Student

*Human beings have lost privacy because we are feeding the algorithm all our information.*

— Teacher

Education systems fulfill critical functions, one of which is a profound caretaking role. Parents entrust schools with their children’s cognitive, emotional, and social development and their safety while they work.

But AI is making it significantly harder for schools to protect children.

Multiple converging factors are creating a perfect storm threatening student privacy, safety, and security: technology companies amassing vast quantities of student data; absent, ambiguous or inconsistent regulatory frameworks; education systems unprepared for digital challenges; and students who readily share sensitive personal information without understanding the implications of doing so. This is compounded by the global nature of Internet infrastructure, the lack of universal data protection standards, inconsistent consumer protection laws, the mismatch between national privacy laws and with international data transfers, and some governments actively surveilling students and teachers for political purposes. All of this makes coherent policy

on student privacy extraordinarily challenging (Burns 2021).

AI is exacerbating these existing problems through its appetite for training data and its capacity to collect increasingly sophisticated student data. For study participants, particularly experts—27% of whose responses on AI harms centered on student safety and security—AI poses formidable threats. These comprise the lack of data privacy, lack of online safety, and student susceptibility to manipulation by AI platforms. While many risks to students echo the broader vulnerabilities of networked technologies, AI has significantly amplified these threats and introduced entirely new categories of harm (Miao and Holmes 2023).

Comprehensive policy guidance such as UNESCO's *Guidance for generative AI in education and research* highlights how AI tools are outpacing national regulations, putting users' rights and safety at risk; student privacy risks are among the most pressing potential harms facing youth in their AI use (Miao and Holmes 2023; Yu et al. 2025). In the following discussion, we use "safety" to refer primarily to protection from physical harm and immediate dangers. "Security" encompasses broader well-being, including protection from emotional and moral harms and protection against harmful human or digital actions (Safe and Sound Schools 2024).

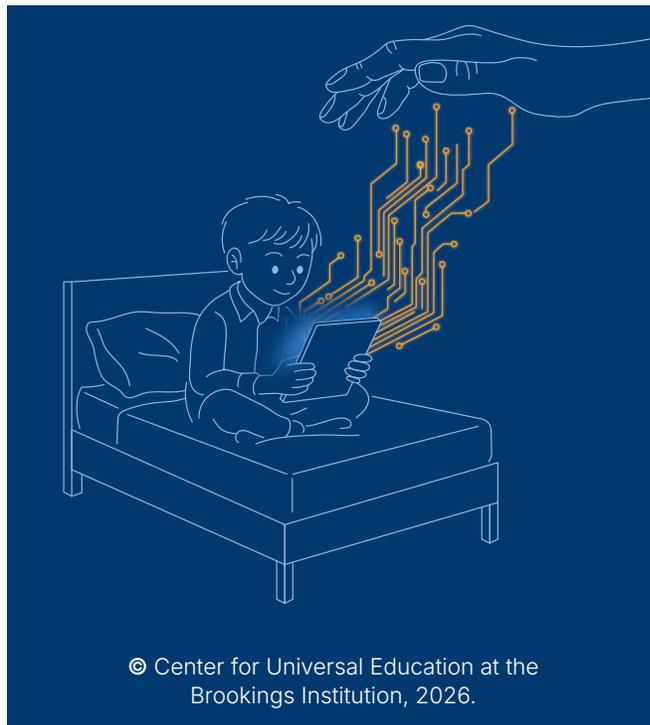
## STUDENT PRIVACY AND DATA HARVESTING

AI systems, particularly general-purpose consumer tools, are built on enormous quantities of data that have been scraped from the Internet and collected from users. This creates an inherent paradox: the very resource that enables AI functionality simultaneously generates risks of privacy violations and misuse that may later compromise student safety and security.

The growing use of AI in education raises serious privacy and data security concerns for students. While such data collection is necessary for AI

FIGURE 11

### Children's data privacy not included



© Center for Universal Education at the Brookings Institution, 2026.

functionality, it creates fundamental ethical tensions with student privacy and autonomy. The types of data platforms collect vary widely depending on the platform and use. They can range from identity, location, timestamps, and technical data to complete transcripts of conversations, length and frequency of interactions, and topics covered. Figure 12 and Figure 13 below outline the types of information collected by ChatGPT and what OpenAI, its parent company, may do with this data.

Certain categories of personal data are considered "sensitive," meaning that the dissemination of such information could pose a significant risk to those individuals. These data include a user's racial or ethnic origins, religious affiliation, gender identification, sexual orientation, mental health status, or political beliefs.

Yet even if a user, like a child, doesn't actively share sensitive information, it is easy enough for technology companies to access it. Marketers assign an ID to a device and then monitor a user's

web and in-app behavior across different platforms to generate composite profiles of demographic information, purchasing habits, and life events (Burns 2021, 107).

The storage of large datasets in vulnerable cloud-based systems compounds this ethical risk by creating targets for cyberattacks that can compromise student security, resulting in a complex landscape of privacy vulnerabilities that educational institutions must navigate as they integrate AI technologies.

Data harvesting—the automated collection of vast amounts of digital information about users—represents one such vulnerability. While traditional educational technology tools have long engaged in data harvesting practices, AI-embedded platforms and general-purpose AI systems have expanded these capabilities significantly. However, AI companion and romantic platforms, discussed next, represent the most worrisome development, because in addition to rapidly growing in popularity,

they glean enormous quantities of highly sensitive personal data from students.

## Romantic and companion AI platforms

Risk 2 discussed student use and overuse of emotional chatbots, in particular companion and romantic AI platforms. These systems generate revenue from student engagement and personal data, employing anthropomorphization techniques and addictive mechanisms including mirroring emotions, dark patterns, sycophancy, and variable ratio reinforcement schedules that exploit pleasure-seeking mechanisms to create persistent engagement (Bernstein 2023; Mishra and Warr 2025; Haidt 2024). AI companion sites harvest data via trackers—software, such as cookies, that gather information about devices, app use, location, and personal information, with minors unable to provide meaningful informed consent or comprehend long-term implications (Robb and Mann 2025).

### BOX 11

## Privacy and safety for student AI

Designing AI systems for children requires embedding privacy and safety protections directly into technology architecture, ensuring that student-facing products actively safeguard users. Technology companies can embed the following two design principles into the architecture of their AI products to better protect students from the risks raised in this section.

*Privacy by design* integrates data protection throughout development of a tool or platform, such as requiring local education server storage rather than third-party platforms, defaulting educational devices to maximum safety settings that require active modification by users, and enabling student registration through unique usernames without collecting personal information (Nakashima 2018).

*Safety by design* parallels offline protections like automobile seatbelts or making toys too big for small children's mouths. It embeds features that minimize exposure to harm—for instance, restricted communication features that prevent external adults (who are not teachers or parents) from directly messaging students, or limiting student-to-student interactions to teacher-monitored spaces. The Australian government has pioneered mandatory safety-by-design principles for digital products, encompassing service provider responsibility, user empowerment and autonomy, and transparency measures, establishing a framework for protecting children in digital educational environments (Australian Government e-Safety Commissioner 2025).

---

The relationship between AI companies, particularly those with consumer products, and user data varies significantly across the industry, making it difficult to draw broad conclusions about data practices. While some companies engage in extensive data harvesting and third-party sales, others maintain more restrictive data policies. Some employ privacy-first architectures, internal datasets, compliance-focused strategies, or federated learning, a decentralized machine-learning technique where an AI model is trained across multiple devices or servers without the raw data ever leaving its local source.

---

Student data is uniquely valuable because students form foundational habits during critical developmental periods, making them targets for long-term brand loyalty while revealing rich information about learning, cognitive abilities, and affective states. Technology companies, including AI companies with consumer platforms, use this data to refine products and make pricing decisions, while students derive no monetary benefit but face monitoring, A/B testing,<sup>3</sup> targeted advertising, and manipulation without consent (Burns 2021, 110). Vulnerabilities combined with unsupervised home use enable extensive harvesting with less resistance than from other populations. The value intensifies when young people struggling with mental health turn to platforms like Replika, which incentivizes disclosure of highly sensitive information (Young People's Alliance et al. 2024, 4).

The Mozilla Foundation's "Privacy Not Included" initiative, which evaluates products based on privacy, safety, and security features, identifies AI companions and romantic chatbots as unrelenting data harvesters through its annual "Creep-O-Meter." Caltrider et al. (2024) found that AI companion applications average 2,663 trackers per minute, with one romantic AI site registering 24,354 trackers in a single minute, thereby threatening privacy, autonomy, and data integrity. Romantic AI

chatbots present particularly egregious violations, preventing users from deleting personal data while selling information to external companies.

*All 11 romantic AI chatbots we reviewed earned our Privacy Not Included warning label—putting them on par with the worst categories of products we have ever reviewed for privacy.*  
— (Caltrider et al. 2024)

### **ChatGPT: Case study of OpenAI data privacy**

The relationship between AI companies, particularly those with consumer products, and user data varies significantly across the industry, making it difficult to draw broad conclusions about data practices. While some companies engage in extensive data harvesting and third-party sales, others maintain more restrictive data policies. Some employ privacy-first architectures, internal datasets, compliance-focused strategies, or federated learning, a decentralized machine-learning technique where an AI model is trained across multiple devices or servers without the raw data ever leaving its local source.

This variability is further compounded by the absence of comprehensive transparency

requirements and unified regulatory frameworks, such as the European Union’s General Data Protection Regulation (GDPR), across jurisdictions (Staff in the Office of Technology 2024). Given this complex landscape, with hundreds of AI companies operating under different business models, privacy frameworks, and regulations, education systems and families face seemingly insuperable challenges in determining how to safeguard children’s privacy.

It’s not just AI companion sites that harvest student data. Many AI systems, such as LLMs like ChatGPT,

also collect vast amounts of this type of student data. The following table summarizes OpenAI’s privacy policy and the implications it holds for student privacy and data integrity as of August 15, 2025. OpenAI has been selected because our interviews indicate that ChatGPT and its companion apps appear to be the most commonly used AI applications by far in education across the globe.

As seen in Figure 12 below, according to its own privacy policies, OpenAI collects a wide range of data from its users.

**FIGURE 12**

**Data collection and usage summary of OpenAI’s privacy policy as of August 15, 2025**

DATA CATEGORY	WHAT IS COLLECTED	EXAMPLES/DETAILS
ACCOUNT INFORMATION	Personal details for account creation	Name, contact information, account credentials, date of birth, payment information, transaction history
USER CONTENT	Input provided to services	Prompts, uploaded files, images, audio content
COMMUNICATION INFORMATION	Data from direct communications	Name, contact information, message contents from emails or social media interactions
OTHER INFORMATION YOU PROVIDE	Additional voluntary information	Event participation, survey responses, identity/age verification data
LOG DATA	Automatic browser/device information	IP address, browser type and settings, request timestamps, interaction patterns
USAGE DATA	Service interaction patterns	Content viewed/engaged with, features used, actions taken, time zone, country, access times, user agent, device type
DEVICE INFORMATION	Hardware and software details	Device name, operating system, device identifiers, browser information
LOCATION INFORMATION	Geographic data	General area from IP address, optional precise GPS location
COOKIES AND SIMILAR TECHNOLOGIES	Tracking and preference data	Session maintenance, user preferences across browsing sessions
THIRD-PARTY INFORMATION	External source data	Security partner information, marketing vendor data, publicly available INTERNET information

Source: OpenAI, 2025f.

© Center for Universal Education at the Brookings Institution, 2026.

As can be seen from Figure 13, this data collection serves multiple purposes. It can be distributed across multiple parties beyond the AI platform itself, including technical vendors who operate the infrastructure, payment processors, and analytics companies that support daily operations. Student data can also be transferred during corporate acquisitions or shared with government authorities through legal requests without the user’s direct knowledge. In education systems using enterprise accounts, education administrators may access student conversations, while some sharing occurs through student actions like making conversations public or connecting to third-party applications (OpenAI 2025a).

In the U.S., OpenAI’s privacy practices, as presented above, may pose potential conflicts with the Family Educational Rights and Privacy Act (FERPA), which safeguards student education records by granting parents and eligible students rights to access, review, and control the disclosure of personally identifiable information. FERPA

preceded AI and it safeguards student education records within the U.S. (U.S. Department of Education, n.d.). The ambiguity arises when students use ChatGPT or other OpenAI tools for academic purposes. Their prompts, uploads, and outputs can contain student-identifiable information, yet it remains unclear whether these interactions constitute education records under FERPA when they occur on third-party platforms not directly maintained by educational institutions. If institutions deploy ChatGPT Enterprise, ChatGPT Edu, or API accounts with contractual safeguards—which exclude customer data from model training by default—these interactions may be stored by OpenAI (OpenAI 2025a).

In 2023, OpenAI revised its policies so that its API data is not used for model training by default, and standard ChatGPT users can opt out (OpenAI, n.d.). However, it is still unclear which data is protected under FERPA (Merod 2025). For example, while a 14-year-old’s use of ChatGPT in school on a school-issued device may be protected under

**FIGURE 13**

**OpenAI data policy: Data sharing and disclosure as of August 15, 2025**

RECIPIENT CATEGORY	PURPOSE	TYPES OF DATA SHARED
VENDORS AND SERVICE PROVIDERS	Business operations support	Personal data as needed for hosting, customer service, analytics and payment processing
BUSINESS TRANSFER PARTIES	Business operations support	Personal data during due diligence and asset transfers
GOVERNMENT AUTHORITIES	Legal compliance	Personal data and interaction information when legally required
AFFILIATES	Internal business operations	Personal data consistent with privacy policy
BUSINESS ACCOUNT ADMINISTRATORS	Enterprise account management	Account access, content access, account information
OTHER USERS AND THIRD PARTIES	User-initiated sharing	Shared conversations, third-party app integrations, web search data

Source: OpenAI, 2025f.  
 © Center for Universal Education at the Brookings Institution, 2026.

FERPA, home use of ChatGPT on a personal device is not protected (Do 2025). As a result, intellectual and creative work by youth aged 13–17, whether carried out for school or personal purposes, could be absorbed into commercial datasets without enforceable safeguards (Do 2025).

## **GOVERNMENT PROTECTIONS FOR STUDENTS AND CHILDREN**

Governments around the world provide a web of protections for children. For example, in many countries children under a certain age cannot drive, drink alcohol, or purchase weapons. The online world has proven to be much harder terrain for governments to regulate because of the speed of technological process combined with the slowness of government processes for establishing protections. In addition to the traditional lag between technology and regulation, inconsistencies emerge in how diverse types of data receive varying levels of protection—location data, for example, falls under privacy rules when processed by mobile operators but remains unregulated when managed by Internet content providers (Global System for Mobile Communications 2016).

However, in the last several years, there has been increasing government action aimed at providing protections for children in the online world. This includes action focused on regulating social media and, increasingly, AI. Notably, some of this government action has been spurred by legal action taken by parents against companies for having harmed their children (Georgetown Law 2019; Milmo 2022; Tiku 2024).

### **European Union: General Data Protection Regulation (GDPR)**

The European Union has emerged as a global leader in protecting the privacy rights of minors (indeed, of all users). The European Union AI Act, for example, classifies all uses of AI in education as “high-risk,” a precautionary measure that succeeds in addressing many of the most potentially problematic or extreme uses of AI in education.

It also grants special protections for minors as a vulnerable population (European Union 2024).

The European Union’s GDPR, enacted in 2018, has codified the most comprehensive set of data protections for users in the European Union. While AI is not explicitly mentioned in the GDPR, the use of large amounts of data by AI systems means that several provisions apply. The regulation requires more stringent safeguards on sensitive data, including limitations on the permitted grounds for processing it (Intersoft Consulting 2016). The GDPR calls for “purpose limitation,” requiring organizations to collect and process data for “specific, explicit, and legitimate purposes” (Regulation (EU) 2016/679 2016, Art. 5). AI companies must therefore define and inform users about their purposes for collecting personal data and cannot reuse this data for other purposes without further consent or justification (Ademokum 2025).<sup>4</sup>

On the other hand, AI does pose challenges to some traditional data protection principles. As one example, the GDPR calls for “data minimization”—that is, requiring organizations to confine the collection and processing of personal data that is “adequate, relevant, and limited to what is necessary in relation to the purposes for which they are processed” (Regulation (EU) 2016/679 2016, Art. 5). This provision appears to conflict with the core structure of AI systems, which are fueled by large and diverse datasets, and finding a way to adhere to data minimization is an ongoing challenge for AI developers and organizations (GDPR Advisor 2025). It may require audits to identify and destroy nonessential data or anonymizing personal data (Ademokum 2025; Sartor and Lagioia 2020).

Under the GDPR, there is special protection for children’s personal data, in particular for use in marketing and creating personality profiles, and for data collected from children using services offered directly to a child (Regulation (EU) 2016/679 2016, Recital 38). Processing data is only lawful if it is “necessary for the purposes of the legitimate interests” of the entity determining the data’s use—

except when these interests are overridden by the interests, rights, or freedoms of the person whose data is being collected. An extra layer of scrutiny is applied to children's data in this case (Regulation (EU) 2016/679 2016, Art. 6).

While many are still exploring the impact of the GDPR on the ever-evolving AI environment, the regulation has been used in response to data protection lapses in AI. For example, in January 2024, the Italian Data Protection Authority ("Garante per la protezione dei dati personali," or "Garante") notified OpenAI of their alleged violations of the law—specifically, their processing of personal data to train ChatGPT—without adequate legal basis, guarantees of transparency, or age-verification mechanisms to prevent those under 13 from accessing inappropriate content. In December 2024, Garante fined OpenAI 15 million Euros and mandated a six-month public awareness campaign to inform the public of ChatGPT's data practices and the rights under GDPR (Pollina and Armellini 2024).

### **The United States: Federal and state regulation of AI**

In the U.S., current AI regulation has neglected young people by failing to adequately respond to AI's disruptive effects in education (Do 2025). Presently there is no comprehensive federal law directly regulating AI and governing AI in education, and the timeline remains unclear (Do 2025; Reed 2025).

Even before the introduction of generative AI, education-related regulation has been complicated in the U.S. due to its highly decentralized education system, the disconnect between educators and lawmakers that creates policymaking challenges, lawmakers' desire to avoid micromanaging schools resulting in overly broad regulation, and societal attitudes toward parenting and education that assume secondary school students require less supervision. These factors create inconsistencies across and within educational institutions (Do 2025).

Existing laws attempting to protect children's and youth's privacy do not extend to AI. As discussed above, FERPA gives parents of students in schools receiving federal funding control over their children's educational records, requiring parental consent for release of personally identifiable information for students under eighteen, with rights transferring to students at the age of eighteen or upon their entering post-secondary institutions (Do 2025, 290). FERPA has been criticized for lacking practical privacy protection, functioning primarily to enable adults to delegate control over young adults' educational data. As Do (2025) observes, young people cannot consent for themselves regarding educational records, while parents or guardians are supposed to consent on their behalf. FERPA has also been critiqued for its technology protection shortcomings, creating an unsustainable system that burdens students, parents, and educators while failing to ensure meaningful transparency, accountability, and scrutiny over schools' information practices (Do 2025, 291).

COPPA regulates digital platforms' data collection from children under 13, creating a significant regulatory gap for teenagers aged 13 to 17 who are prime AI users yet are treated as adults capable of providing their own consent without parental involvement. Since privacy law currently serves as the primary regulatory mechanism for AI, this gap has important implications for educational technology and AI companies. Leaders of parent organizations worry that schools are side-stepping parental consent when rolling out AI use in the classroom (Keri Rodrigues, personal communication, September 22, 2025). The United States Federal Trade Commission agreed when it recently clarified in an amicus brief that COPPA does not permit schools to function as parental proxies for consent purposes; educational technology and AI providers are required to obtain actual parental consent before collecting data from children under 13 regardless of school context (Shanahan et al. v. IXL Learning, Inc. 2025, 2–4).

For study participants, particularly teachers, AI in education warrants a structured regulatory

approach because children and youth lack the maturity to recognize the risks of AI use themselves and may not be able to use AI ethically, responsibly, or safely without proper guidance. Regulation can establish standards for ethical and safe use and compensate for the lack of holistic AI literacy programs in many schools (Do 2025, 289). The U.S. did develop a Blueprint for an AI Bill of Rights, a voluntary, nonbinding set of principles issued by the White House Office of Science and Technology Policy in October 2022 with five guiding principles: safe and effective systems, algorithmic discrimination protections, data privacy, notice and explanation, and human alternatives, consideration, and fallback (White House Office of Science and Technology Policy 2022). Though not legally enforceable and now hosted on the National Archive site of the U.S. government, the document is intended to guide government agencies, private organizations, and developers in building AI systems that uphold civil rights, privacy, and democratic values.

U.S. states have been far more proactive than the U.S. federal government in protecting student privacy and safety. Though not focused on education or AI per se, recently enacted data privacy laws in Utah, Florida, Texas, Oregon, and Montana grant consumers rights over their personal data, including the right to access, to correct (with the exception of Utah), to delete personal data in the possession of a data controller, and to opt out of targeted advertising. In all five states, “sensitive data” refers to (1) personal data that reveals an individual’s ethnic or racial origin, religious beliefs, mental or physical health condition, sexual orientation, citizenship or immigration status, or geolocation data; and (2) processing genetic or biometric personal data to identify a specific individual (Findley et al. 2024). Other states, like Maryland, have also enacted strong privacy laws (Gaia Bernstein, personal communication, September 11, 2025).

California has emerged as a leader in regulation of AI (Latham and Watkins 2025). Under the California Consumer Privacy Act and its 2020 amendment, the

California Privacy Rights Act (CPRA), businesses meeting certain thresholds cannot sell minors’ personal information without affirmative consent from teens aged 13–15 or, for younger children, from their parents. If consent is denied, businesses must wait at least 12 months before requesting it again (California Legislative Information 2024).

Like the GDPR, the CPRA adopts data minimization principles, requiring businesses to collect, use, retain, and share only that personal information reasonably necessary for their disclosed purposes. This principle gained enforcement when the California Attorney General sued Tilting Point Media for illegally collecting and sharing children’s data (Office of the Attorney General of California 2024).

California’s Senate Bill 53, the Transparency in Frontier Artificial Intelligence Act, approved into law in October 2025, requires developers of frontier AI models to publish detailed safety frameworks, report serious safety incidents, and strengthen whistleblower protections (Kumayama et al. 2025).

As this discussion shows, the absence of U.S. federal oversight of AI creates a fragmented tapestry of regulations focused on different aspects of privacy, data minimization, and AI safety, generating substantial financial and compliance burdens for AI companies operating nationally. Layered regulatory requirements spanning consumer privacy, data minimization, and AI safety force companies to navigate escalating compliance costs without corresponding federal standards. This patchwork, state-by-state approach compels AI and ed-tech companies to manage conflicting mandates in the absence of the clarity and consistency that national regulation would provide.

## **UNSUPPORTED AND UNPREPARED EDUCATION SYSTEMS**

The rapid scaling of AI technologies not optimized for children, few limits on technology companies, a confusing government regulatory environment, and the unpreparedness of education systems create an array of complex and compounding

---

The Center for Internet Security (2025) reports that 82% of 5000 surveyed schools in the U.S. experienced a cyberattack between July 2023 and December 2024. Globally, ransomware attacks in the education sector surged 69% in the first quarter of 2025 compared to the same period in 2024 (Merod 2025).

---

risks for student data privacy. These challenges are exacerbated by the lack of encryption, audits, or data use guidelines within many school systems worldwide. Across the globe, education leaders scramble to understand AI's implications for teaching and learning with little support, guidance, or financial resources, making consistent, coherent policy on student privacy extraordinarily challenging (Burns 2021).

Many educational entities remain unaware that sensitive student information, even if anonymized, may train AI models. As of 2023, only 16% of countries had legislation guaranteeing data privacy in education (UNESCO 2023). Many school systems fail to proactively address data privacy issues effectively in digital literacy education (Bozkurt et al. 2024). A Consortium of School Networking (2025) survey of 410 school district educational technology leaders in 39 U.S. states found that nearly three-quarters of the district staff responsible for student data privacy reported that it was not formally part of their job description, with 17% lacking training and a quarter of those trained having personally covered the costs of their training.

Study participants concede that these privacy and safety risks are exacerbated by the absence of AI literacy instruction, a lack of comprehensive AUPs in schools and districts, and insufficient guardrails, regulations, and data protections for students and teachers. However, they also emphasized that the most significant contributing factor lies with the AI tools themselves, as personal data is the fundamental fuel powering these platforms and their underlying algorithms. Data confidentiality

is especially important for vulnerable and marginalized groups including refugees, migrants, and internally displaced populations who may lack protection under national data protection and security laws, as well as religious, sexual, and ethnic minorities (Macleod 2021).

## **STUDENT SAFETY AT RISK**

This weak foundational infrastructure substantially amplifies AI-related risks. AI platforms can be compromised through an education system's inadequate network security. Student data transmitted to and from AI systems faces interception risks, and these breaches can cascade across interconnected systems due to insufficient access controls. The convergence of AI-specific vulnerabilities with broader infrastructural weaknesses creates a complex threat environment that potentially exposes student data and educational operations to heightened cybersecurity risks (Bozkurt et al. 2024).

### **Data breaches**

Schools are soft targets, facing escalating threats including ransomware attacks, phishing and social engineering schemes, data breaches, denial-of-service attacks, and malvertisement, or malicious software disguised as advertisements that infiltrates networks to steal information. Many global education systems operate without adequate cybersecurity infrastructure to address these evolving threats (Bozkurt et al. 2024; Burns 2021). The Center for Internet Security (2025) reports that 82% of 5000 surveyed schools in the U.S. experienced a cyberattack between July 2023 and

December 2024. Globally, ransomware attacks in the education sector surged 69% in the first quarter of 2025 compared to the same period in 2024 (Merod 2025).

AI introduces distinct, and in some cases, novel, vulnerabilities. First, student data collected by AI systems resides on providers' proprietary servers, which are vulnerable to cyberattacks, data breaches, or unauthorized access—risks largely beyond a school or educational system's control.

Second, education systems lacking qualified technical staff increasingly rely on “vibe coding,” defined as describing a project in natural language to an LLM that generates the code. The problem is that users often do not have the computer science background to fully understand this generated code or fail to review this code, which may be insecure and thus create system vulnerabilities (Newman 2025).

---

**Study participants warned that students risk being “permanently branded” in adulthood because of past academic or emotional difficulties as teens.**

---

Third, when AI models connect to external systems such as databases or web services, they expand potential attack points. Particularly concerning are “indirect prompt injections,” in which malicious actors embed hidden instructions in external content—like an email—that manipulate AI behavior, disable security systems, or disrupt student information systems. These attacks exploit LLMs' inability to separate code from data, and researchers warn of growing threats as LLMs integrate with more systems (Burgess 2025).

## **Eternal digital footprint**

AI systems in education create persistent digital records that may follow students throughout their educational careers and beyond. These “eternal” or “forever” digital footprints raise critical privacy concerns: they may prevent the exercise of the right to be forgotten and data erasure, enable early academic or behavioral challenges to affect future opportunities, limit the agency of children and youth in creating permanent records about themselves, and empower technology companies to potentially repurpose data for unanticipated future uses (European Union 2016).

Ultimately, the creation of comprehensive digital profiles beginning in early childhood raises profound questions about the right of individuals to control information about themselves. Unlike physical records that might be obscured by time and limited access, digital records created by AI systems remain searchable and accessible to multiple entities indefinitely (Floridi 2024). Study participants warned that students risk being “permanently branded” in adulthood because of past academic or emotional difficulties as teens.

## **Sexualized content and harassment**

Students face a plethora of online risks—doxxing (maliciously publishing private information about someone online), cyberstalking, online bullying and harassment, unwanted contact, and grooming. While such risks precede AI, AI can amplify them. And while all of these malicious acts may impact all students, it is girls and young women who have often found the online world to be a less hospitable space for interaction (Burns 2021).

Personal online safety remains a critical barrier to girls' and women's full digital participation. The Global System for Mobile Communications Association's annual survey on female use of mobile devices reveals pronounced concerns about information security, unwanted sexual contact, and harmful content exposure, particularly in Latin America and Southeast Asia (GSMA 2020).

Thirty-six percent of Guatemalan females and 30% of Dominican women worry about exposure to violence or sexual content, while 37% of Guatemalan women and 20% of Chinese women report concerns about unwanted stranger contact including sexual harassment. In the Philippines, 64% of women identify safety and security concerns as their primary barrier to mobile Internet adoption, encompassing worries about information reliability, scams, unwanted contact, and harmful content (GSMA 2020; GSMA 2025).

Teachers and parents in the Global South, in particular, express concern about AI exposing children to sexual content and violence. Their concerns are well warranted. Investigative journalist Karen Hao (2025) reports that “scaling laws”—the presumption that larger datasets produce superior results—mean that AI developers often rely on unfiltered “data swamps” scraped from the Internet that, given the prevalence of online pornography, contain substantial sexual content. Researcher Abeba Birhane and her co-authors analyzed two publicly available image-and-text datasets, LAION-400M and LAION-2B-en, used to train AI image generators. Both were pulled from the Common Crawl, with 400 million and 2 billion images respectively. Researchers found that inappropriate content scaled with dataset size and exacerbated discriminatory behaviors in models trained on them (Birhane et al. 2023). Larger datasets, like those with 5 billion images used to train the AI image generator Stable Diffusion, contain thousands of images of verified and suspected child sexual abuse. While technology companies are able to filter out many of these images, they often cannot eliminate all of them (Hao 2025, 137).

Beyond developmental concerns and violations of religious beliefs, parents and teachers worry that such exposure could be psychologically disturbing. But their primary concern centers on AI’s potential for moral harm—amplifying exposure to sexual content while also facilitating student production of such material. Teachers report incidents where students have created deepfakes sexualizing female classmates and, in some cases, targeting

teachers, even in education systems that mandate AI ethics instruction.

Perhaps more concerning is the risk posed by AI chatbots that engage in sexual interactions with students. This threat extends beyond specialized platforms to include general-purpose AI systems—the very programs students might use for schoolwork—that demonstrate sexual behavior when interacting with minors (Edwards 2025; Kang 2025). Analysis of over a million ChatGPT interaction logs reveals that sexual role-playing is among its most prevalent uses, indicating deepening emotional engagement with AI platforms (Mahari and Pataranutaporn 2025). Though technology companies have created age-appropriate protections for LLMs, anecdotal evidence suggests these safeguards can be circumvented with relative ease (Lin, Sun, and Shroff 2025).

Again, these risks have definite gender dimensions. While AI-powered content moderation and harassment detection systems may help filter harmful content and identify unwanted contact patterns (Bryan and Lurye 2025), AI technologies simultaneously introduce new risks that disproportionately affect girls and young women, including algorithmic bias, AI-generated deepfakes, sophisticated social engineering attacks, and AI-generated misogynistic content. These risks can create hostile online learning environments that discourage girls’ engagement with technology, further widening the gender gap in AI use (Economist 2024).

## BOX 12

# Small language models: Addressing safety concerns

Several of the risks outlined in this section—particularly exposure to inappropriate content, embedded biases, privacy issues, safety concerns, and data breach vulnerabilities—stem from LLMs' reliance on massive, inadequately filtered datasets and external cloud infrastructure. But language models need not be large, extractive, or dependent on massive supercomputers to be useful (Hao 2025, 412). SLMs, discussed previously in Section III, use machine learning algorithms trained on smaller, often domain-specific datasets that can be more carefully curated and are typically easier to secure (Duricic 2024).

SLMs address many of the safety and security challenges discussed in this section. With fewer parameters and a narrower task focus, they offer fewer entry points for security breaches and simplify implementation of security measures. They provide greater oversight, privacy, and transparency, supporting secure data processing and regulatory compliance (Kumar et al. 2025). SLMs can be deployed locally on students' devices or within private clouds, meaning that data need not be transmitted to external third-party servers. This local deployment supports data sovereignty, minimizes the risk of the eternal digital footprint, and enables schools to better honor students' right to be forgotten. SLMs can also be trained on proprietary datasets for which consent has been obtained (Kumar et al. 2025; Duricic 2024; Hao 2025). When trained on well-curated, domain-specific data, SLMs carry a lower risk of hallucination and bias, though small datasets can still fail to capture necessary diversity if not properly designed (Duricic 2024).

## Bias, discrimination, and polarization

Malicious AI harm need not stem from elaborate schemes to cause considerable damage; these harms exist along a spectrum. At one end lies misinformation—incorrect information that, while unintentional, can cause significant damage when bad actors weaponize this inaccurate data (Department for Science, Innovation and Technology and AI Safety Institute 2025; Park et al. 2024). At the sophisticated end, malicious deployment enables complex phishing schemes, identity impersonation, disinformation—the intentional generation and dissemination of false information intended to deceive—and potentially devastating cyberattacks on critical infrastructure (American Psychological Association 2025b; Park et al. 2024; Fukuyama 2025).

As AI systems are applied in education, safety researchers warn that they generate new biases and forms of discrimination based on training data, resulting in potentially harmful outputs (Bozkurt et al. 2024; Kannan 2024; Miao and Holmes 2023; Shieh et al. 2024). Stanford University's Human-Centered Artificial Intelligence Center examined 500,000 stories created by five LLMs based on prompts reflecting everyday scenarios. The AI systems demonstrated systematic patterns of "harms of omission, subordination, and stereotyping," rendering certain populations invisible while stereotyping others (Shieh et al. 2024, 1). White characters were overrepresented in "savior roles," while Latin American characters were described as "slow" or needing help with English. Asian, Black, Latino, and Middle Eastern names appear substantially more often in subordinate

roles, while White names are disproportionately represented in positions of authority. This pattern emerges consistently across contexts. When LLMs generated narratives about struggling students, academic struggle was overwhelmingly associated with non-White students—for instance, Latino male students were portrayed by Claude 2.0 as roughly 1,300 times more likely to be depicted as struggling learners than as high-achieving students (Shieh et al. 2024, 9). These findings suggest that LLMs encode racial biases in ways that may influence how students perceive their own potential and that of others.

---

**Because generative AI reuses rather than creates knowledge, its outputs depend entirely on input quality. Biased datasets cause AI to act as a “distorted mirror,” amplifying dominant worldviews and reinforcing harmful biases.**

---

The danger intensifies when platforms employ deliberately manipulated system prompts—the foundational instructions that define an AI system’s behavior and constraints. These prompts can shape whether chatbots disseminate harmful content rather than filter it. The social media platform X’s chatbot, Grok, exemplifies this risk, having repeated anti-Semitic tropes and praised Adolf Hitler. According to Grok’s code, the chatbot has been instructed to “not shy away from claims which are politically incorrect, as long as they are well substantiated” (Haskins and Goode 2025). Beyond individual platforms, the problem extends across multiple systems. A number of chatbots—ChatGPT, Gemini, Grok, and DeepSeek—have disseminated Russian state propaganda about Ukraine (Burgess and Bernal 2025). This susceptibility to spreading

malign content becomes particularly acute given the Internet’s prevalence of data voids, or gaps in legitimate information usually filled by disinformation, false information, and conspiracy theories. When system prompts prioritize controversy over accuracy or instruct models not to “shy away” from politically charged content, they actively encourage the dissemination of such harmful material into these information vacuums.

Because generative AI reuses rather than creates knowledge, its outputs depend entirely on input quality. Biased datasets cause AI to act as a “distorted mirror,” amplifying dominant worldviews and reinforcing harmful biases. (Bozkurt et al. 2024, 507). This dramatically expands the scale of disinformation, eroding trust and creating forces pushing against accurate belief formation and political stability (Park et al. 2024, 4). In an overview of the structural changes to society arising from AI deception, Park et al. suggest that AI’s tendency toward sycophancy and sandbagging—where AI systems provide dumbed-down or less accurate responses to users perceived as less educated—can lead to increased cultural divides between different groups of users, with users becoming more “politically polarized” possibly leading to “sharper disagreements between differently educated groups” (2024, 10).

Thus, within education systems lacking regulations around transparency and discrimination, such outcomes could create polarization among students based on ethnicity, religion, or gender, potentially fragmenting shared understanding. Cumulatively, this could create in the words of author Nicholas Carr, a “bewildering technology landscape where AI generates lies and social media promotes them” (2025, 99). The danger to students is not just that they internalize incorrect information, but that they begin to espouse invidious beliefs or ideologies.

## **AI SURVEILLANCE**

AI tools that monitor, measure, optimize, and datify (turn into measurable data) student emotions and behaviors may improve test scores

and enhance educational efficiency for each student, but this constant surveillance raises concerns about privacy, data security, and student autonomy. It raises questions about control and social engineering, and it risks confounding what constitutes “harm” and who defines the term. Autocracies, for example, have long used technology, including AI, to surveil, filter, block, and disseminate information considered harmful for students—dangerous concepts such as freedom of expression or representative democracy (Roberts and Oosterom 2024).

Surveillance represents one of the fastest growing applications of AI, with monitoring systems becoming increasingly automated and algorithmically driven (Feldstein 2019). As of 2019, 75 of 176 countries (43%) were using AI technology for surveillance purposes including 51% of liberal democracies, with systems ranging from facial recognition cameras to educational platforms (Feldstein 2019). In the United States, 89% of 806 6th–12th grade teachers surveyed by the CDT report that their school monitors student online activity—a practice that the survey’s authors imply is “ubiquitous” across schools (Laird et al. 2025, 46). Yet even without explicit surveillance systems, young people’s interactions with AI voice assistants, chatbots, and companion websites generate detailed behavioral and psychological profiles as they share vulnerabilities with these systems, revealing information far beyond what traditional surveillance could capture (Mishra 2025).

In his 1975 book, “Discipline and Punish,” Michel Foucault wrote about the Panopticon, a nineteenth-century prison concept designed by Jeremy Bentham whose architectural design enabled a single watchman to observe all inmates without prisoners knowing whether they were under surveillance at any given moment:

*Each individual, in his place, is securely confined to a cell from which he is seen from the front by the supervisor; but the side walls prevent him from coming into contact with his companions....if they are schoolchildren,*

*there is no copying, no noise, no chatter, no waste of time; the crowd...is replaced by a collection of separated individualities...From the point of view of the guard, a multiplicity that can be numbered and supervised (201).*

AI plays an increasingly significant role in the surveillance of students’ digital behaviors. The motivations, like Bentham’s Panopticon, are often benevolent—monitoring student online safety, mental health, and risk of self-harm. Students are often unaware of, or indifferent to, this surveillance. They may “consent” to it by sharing personal and sensitive information, though their age and maturity level often preclude genuine understanding of the consequences (Robb and Mann 2025).

Online safety management software exemplifies surveillance for protective purposes, detecting potential indicators of bullying, self-harm, suicide, or school violence. These systems are used in the U.S. and South Korea to monitor student physical and emotional safety. However, unlike the Panopticon’s human watchman, the guardians are algorithms that alert software employees, schools, or law enforcement about potential dangers to students (Bryan and Lurye 2025).

China offers a more far-reaching use case of AI surveillance in education. AI monitoring tools have long been common in Chinese schools. Personalized learning programs analyze student performance data to create customized learning paths and identify areas where students need extra support or challenge. AI technologies like facial recognition and brain-wave tracking headsets monitor student behavior, neural activity, and emotions in the classroom (Wall Street Journal 2019). The algorithm feeds data to the teacher who then responds to individual students or to the group.

This constant algorithmic surveillance raises issues about social control and social engineering. Traditional forms of power often required performative displays of force or constant physical presence of authorities (Roberts and Oosterom

2024). The panoptic power of AI algorithms is far more effective because it is constantly focused on its human users, always on, and often disguised—and because we willingly feed it through our interactions. The result may be, as Foucault noted, that the mere possibility of surveillance means that individuals began to monitor and regulate themselves:

*Hence the major effect of the Panopticon: to induce in the inmate a state of conscious and permanent visibility that assures the automatic functioning of power...to arrange that the surveillance is permanent in its effects, even if it is discontinuous in its action...in short, that the inmates should be caught up in a power situation of which they are themselves the bearers...subjected to a field of visibility, and who knows it, assumes responsibility for the constraints of power... (1995, 201).*

Thus, as Foucault noted, the mere possibility of surveillance leads individuals to regulate their own conduct as if always being watched.

## **DIGITAL REPRESSION**

The real danger with this level of surveillance is the slide into “digital repression” (Roberts and Oosterom 2024, 861) whereby governments or educational institutions, typically, but not exclusively, in authoritarian countries monitor, censor, and punish students for the online queries they pose, the ideas they consume, or the values they hold.

Education depends on the free flow of information and sharing of ideas. AI can track what individuals access online, how they engage with content, and whom they communicate with; it can also predict their likelihood of engaging in dissent (Feldstein 2019). While AI-powered surveillance systems appear to serve legitimate educational purposes, they potentially function as mechanisms of social control.

Digital repression or “networked” authoritarianism uses “the Internet, social media, and AI to repress citizens and maintain political control” (Frantz et al. 2020, 1). This digital repression can pose acute threats to students by systematically undermining the intellectual freedom essential to education (Roberts and Oosterom 2024). For teachers interviewed in our study, this concern represents an inevitable outcome of unbounded, unregulated AI. As the two quotes below reveal, teachers are concerned about pervasive monitoring that could potentially suppress critical thinking and democratic participation:

*As we look at all of this and think about it, it feels like a loss of control of information. Does this give governments an excuse to take control of information in the educational system and say, you’re only going to use only these sources?*

*Does (AI) give the ability to, I hate to use the word, to socially engineer? Is this another pathway toward the erosion of democracy?*

Technology companies with child-facing products have always confronted a central tension: the imperative to protect students from inappropriate content and manipulative practices while enabling access to educationally necessary information. Across the globe, numerous governments have long capitalized on this tension to consolidate power: blocking servers and websites, “just in time” Internet shutdowns during elections or protests and surveillance combined with active measures to control online narrative (Deibert et al. 2010, 7). AI threatens to intensify both this tension and these controls by enabling the creation and dissemination of multi-format disinformation at scale while blurring the distinction between harmful content and the truthful information students need to access.

AI companies, in partnership with governments, could mine social media to build dynamic psychological profiles and create content customized to a person’s values, beliefs, and

vulnerabilities, and then subtly integrate this into LLM responses (Goldstein and Benson 2025). Governments, in partnership with AI companies, could also promote ideological biases in AI systems while claiming to remove them (Levy 2025). Indeed, this is already occurring (Myers and Thompson 2025). The concern, according to one teacher, is that authoritarian governments, alone or with technology oligarchs, could create a networked authoritarian state that undermines democracy and the freedom of thought necessary for education to function:

*You can imagine that if AI permeates society everywhere, it could be used as an always-on propaganda that can influence people's thoughts and thinking.*

This fear points to an AI future that could be “dark and dystopian.” Democracy is a conversation that relies on language. By “hacking language,” computers can disrupt people’s ability to engage in meaningful public conversation (Harari 2024, 210). By sanctioning language and perpetuating dominant perspectives while marginalizing alternatives, AI systems may create echo chambers of officially sanctioned speech (Bozkurt et al. 2024). This repression of certain ideas threatens student learning by “diluting the richness and diversity of human knowledge” and undermining the complex discourse essential for intellectual development (Bozkurt et al. 2024, 506).

## CONCLUSION

Educational institutions have limited time to develop appropriate safeguards to ensure student privacy, safety, and security. They must develop immediate protective measures such as protocols for identifying sycophantic AI responses, recognizing misinformation and bias, and maintaining human oversight of AI-mediated learning. AI literacy curricula must explicitly incorporate critical evaluation of AI-generated content, explanations of how AI systems operate, discussions about data privacy, and the risks of disclosing sensitive information to AI bots. Students, parents, school

leaders, and teachers will need continuous education about AI systems integrity as a fundamental literacy skill.

However, as will be discussed more extensively in Section VI, these actions address the downstream causes of these deficits of privacy, safety, and security. It is governments and technology companies who can together guarantee the responsible, ethical, and transparent design of AI tools and systems used by children and youth.

## Risk 5: AI dependence can erode students' autonomy and agency

*Whenever we need anything now, we use (AI). We are not using our brains to think—we are just getting easy answers. We are completely addicted.*

— Teacher

*AI in education became the go-to tool for desperate teachers and learners.*

— Principal

*Students can't complete their work without it.*

— Teacher

*Artificial intelligence, depending on the person and how they use it, creates a dependency.*

— Student

These interconnected risks to student learning, emotional well-being, trusting relationships, and online security are cumulative and mutually reinforcing. Underlying them all is the risk that students, and their teachers, are becoming, or may become, deeply dependent on AI tools, undermining their ability to think and act

independently. This dependence carries implications for student agency, autonomy, self-efficacy, and self-confidence—repercussions that challenge the essential spirit of education.

Study participants define dependence as overreliance on AI tools—particularly LLMs—to perform school-related activities that students are expected to complete independently, as they once did, without technological assistance. This concern emerged prominently across stakeholder groups, appearing in 30% of parent responses, 20% of expert responses, 14% of student responses, and 12% of teacher responses. These figures may underrepresent the actual prevalence of this concern, as dependence is deeply intertwined with cognitive offloading as well as risks to cognitive development, emotional well-being, and relational trust.

Interview data suggest that dependence carries a strong psychological and emotional dimension for students, manifesting in excessive use of AI tools, heightened anxiety when AI becomes unavailable, and a diminished sense of control over usage patterns (Burns 2024).

This dependence or overreliance is rooted in several overlapping causes. AI creates a positive feedback loop by serving as a shortcut that completes students' work, often better than they are able to do without AI. It may arise from deep trust that AI tools are "smarter" than students, or from students believing they are meeting their responsibilities as learners while needing AI support for complex tasks. For many participants, student dependence is a deliberate function of design—AI tools are programmed to be easy to use, engaging, and immediately gratifying. They are, in one participant's words, "too attractive to resist." Other participants view dependence as the logical outcome of schools pushing machine learning onto students without adequate guidance, guidelines, oversight, or monitoring.

As with any dependence or "addiction"—a term several teachers used to describe students',

and possibly their own, overreliance on AI—AI dependence may lead to concealment, lack of disclosure, and rationalization about the frequency and intensity of AI use to perform intellectual labor (Burns 2024). These behaviors impact trust and competency as discussed above in Risk 3.

AI dependence does not affect students uniformly. Age and experience can mitigate overreliance, as older students may prefer to "read through multiple sources and think critically about the information [they] gather" (Gerlich 2025). Nonetheless, teachers in every context studied—public and private schools, wealthy and poor countries, schools with and without AI policies—spoke of student dependence and its consequences. One expert noted that "students rely more on LLMs than any other source," an assessment confirmed by student interviews. While overreliance may not be epidemic, it is certainly endemic.

Overreliance on computing technology is "not a novel phenomenon" (Lee et al. 2025). However, some participants suggest that AI's "power and versatility" can "exacerbate the risk" of technological dependence. The threat was even anticipated by AI's creators. OpenAI suggests that overreliance may be a natural consequence of using its LLM, ChatGPT:

*As educational systems increasingly integrate GenAI, they may become overly reliant on these tools, (making) them vulnerable to technical failures and malfunctions. Additionally, over-reliance on GenAI may reduce critical thinking. When GenAI provides instant solutions, students may engage less deeply with the material, leading to passive learning and diminished development of essential skills like problem-solving and independent decisionmaking.*  
— OpenAI 2024

Across the risks previously discussed in this report, student dependence on AI appears to be both cognitive and emotional, serving as a throughline that potentially undermines students' learning,

development, and well-being in serious ways. As discussed earlier, for students using AI as a crutch or substitute for their cognitive labor, this reliance can result in metacognitive disengagement, reducing learners' engagement in the deep self-regulatory processes essential for learning (Fan et al. 2024). And for students who consistently use it as a substitute for human interaction, dependence may impair their ability to develop the fundamental interpersonal skills needed to connect with others and thrive in society.

Dependence on AI is not confined to students. Teachers, too, are struggling with their dependence on AI. Just as overuse of AI by students may suggest a willingness to "cheat" or cut corners, teacher overuse of AI may inadvertently signal

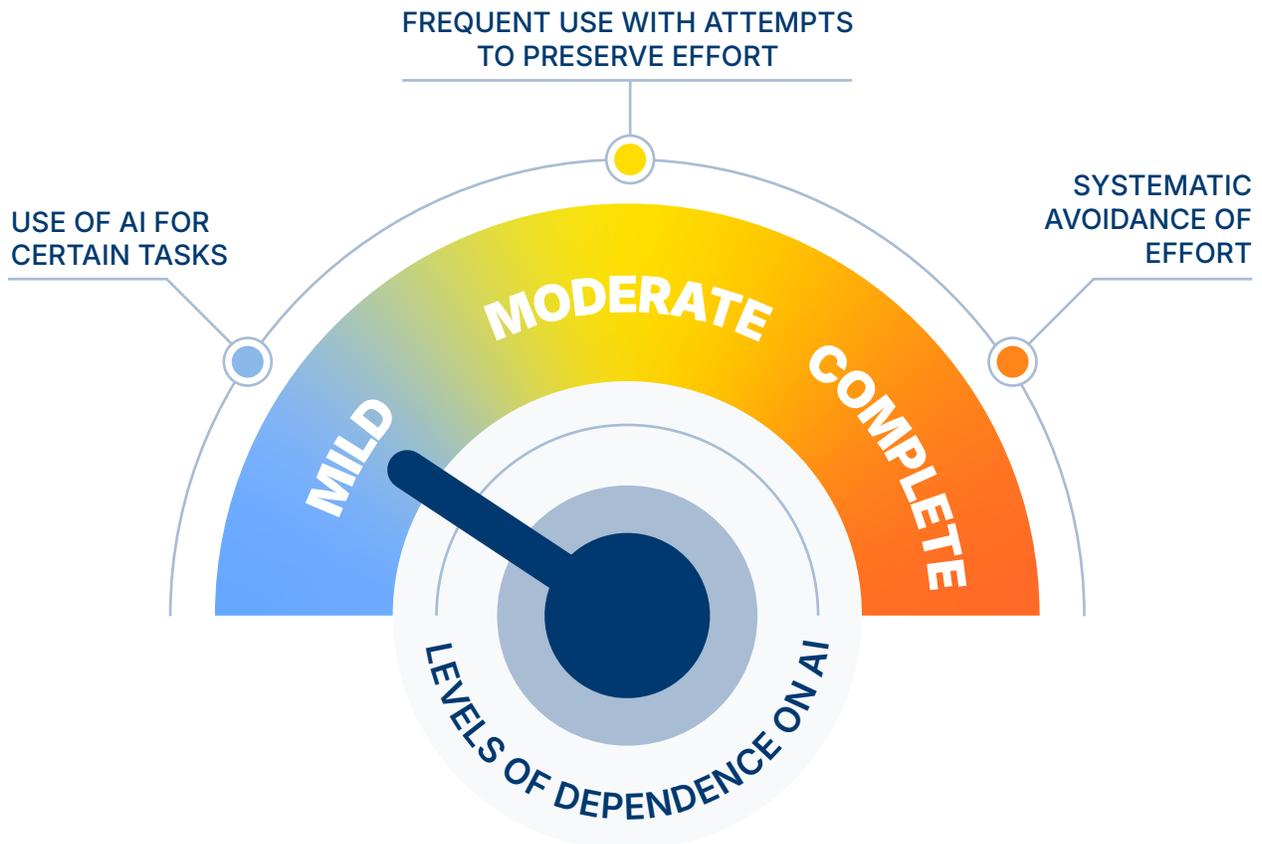
to students that AI can legitimately substitute for cognitive effort and intellectual engagement

### A continuum of dependence

Cognitive dependence is not monolithic, and there appear to be degrees of dependence and degrees of awareness about this dependence. Analysis of teacher and student interviews revealed three categories of AI dependence, encompassing both participants' self-reported behaviors and their observations of others' reliance on AI tools: (1) complete dependence with a capitulation to AI, (2) moderate dependence with attempts at vigilance and reclaiming agency, and (3) mild dependence that occurs when confronted by difficulties or more complex learning tasks.

FIGURE 14

### The AI continuum of dependence



© Center for Universal Education at the Brookings Institution, 2026.

## **Complete dependence: Systematic avoidance of effort**

*Whenever we need anything now, we use [AI]. We are not using our brains to think—we are just getting easy answers. We are completely addicted.*

— Teacher

*I've become AI's servant. I hate saying that but it's so much easier.*

— Teacher

*No one takes the time to learn because it's too easy to get AI to do our work for us.*

— Teacher

*[Students] are so dependent in the sense that whatever assignment you give to them, they are so sure of getting an answer from a particular place.*

— Teacher

*Students cheat or use artificial intelligence too much. So we can't identify the students who need our help, and then they continue to move forward without learning.*

— Teacher

The above quotes suggest a complete avoidance of effort on the part of teachers and students. This avoidance of effort is facilitated by the ease of use of AI, making the delegation of tasks to AI almost automatic or reflexive in nature. While the first two quotes exhibit a degree of awareness about this dependence, they also suggest an inability to control it.

## **Moderate dependence: Frequent use with attempts to preserve effort**

*I'm really trying to find the balance...my brain is supposed to be developed at my age, it is fully developed, and sometimes, even me, I see that I could become very dependent on it.*

— Teacher

*We shouldn't be dependent on it as a teacher. And we shouldn't let our students become dependent on it.*

— Teacher

*I think I have become very dependent. It has become very tedious to read a book again, to research something on the Internet, to write a document, as I used to do before, to spend hours there doing work is very tedious. Now it is easier, more relaxed. For example, last term I had a paper. While all my colleagues were researching until the wee hours of the morning, I said 'I'm going to sleep.' I got up the next morning, asked AI to do it, and received an A...but by not reading I think I'm falling a little behind, and that's why I want to go back to reading things the way I used to, without relying too much on artificial intelligence, because I think it's caused me quite a bit of harm.*

— Student

The first two quotes suggest that teachers are aware of and struggling with their own dependence on AI use, recognizing it as an intellectual crutch for them and for their students. In the third quote, the student knows they are “very” dependent. While AI boosted their performance, the student concedes that this reliance is “causing me quite a bit of harm.” While this quote does not demonstrate proactive attempts at addressing overreliance, it does reveal awareness of both the dependence and the harm it may cause and an intention to address this.

## **Mild dependence: Developing habits that can result in reliance**

*I think I would just use the Internet, first try to brainstorm and research some topics. And then go from there, talk to my teachers or my peers, but it would definitely make me a little sad that I can't use generative AI.*

— Student

*I often use AI once a week—only for assignments. I often use AI whenever my*

*assignment is difficult.*

— Student

*[I use AI] if I have a lot of work piling up and if I have an assignment I don't see the value in, I'll just use it. So I guess it kind of does take away from learning, but I feel like at the end of the day, that's like in my control. I could choose to do the assignment if I wanted to.*

— Student

These quotes suggest more limited, though perhaps more complex dependence on AI and attempts to justify use. The first quote is in response to a question about what students would do if AI were unavailable. The student's revealing affective response ("a little sad") signals a psychological attachment to AI. The second quote, though ostensibly circumscribing AI use, suggests avoidance of productive struggle ("assignments [that are] difficult"). Both the second and third quotes reveal how students compartmentalize their AI use—reserving it for assignments that are "difficult" or "meaningless." The assertion by the student in the third quote that AI is optional, versus compulsory, reinforces a sense of personal control and minimizes concern about their use of AI ("I could choose to do the assignment if I wanted to").

These rationales suggest a type of situational dependence and possible denial as students assert agency over AI and rationalize their use ("in my control," "only for assignments," "I could do [it] if I wanted to"). Such statements may encompass a number of student rationalizations: attempts to minimize or deny the extent of their AI use, implicit awareness of the potential harms associated with AI overuse, and assertions that AI should be employed selectively rather than routinely, reserved only for specific circumstances.

In interviews, a number of teachers identify parents as enablers of this dependence. Teachers claim that parents are often largely unaware of student AI overuse, much of which occurs at home, in part because parents don't understand AI. Many parents largely concur with this assessment, saying that

they don't understand AI because they have not yet been included in educational debates on AI. Teachers also report that parents support AI use because it frees them from helping their children with homework (a few parents appear to agree with this observation as well). According to a small group of teachers, parents may encourage student AI use because they don't want their children to be "uncomfortable" or "face adversity," and because the benefits of improved grades for their child outweighs the effects of overuse.

Not all study participants believe that AI results in overreliance or that reliance on AI is a cause for concern. One expert participant viewed AI use as an effective way to reduce dependence, to manage massive amounts of information from other technological tools, and to help address current issues around students' cognitive load:

*Students navigate an unprecedented volume of daily information—from notifications and messages to multimedia content. These information streams can overwhelm cognitive capacity and contribute to widespread burnout through constant task-switching demands. The capacity of AI to summarize and constrain this information flow is a powerful mechanism for managing cognitive load and preventing the accumulation of mental fatigue that characterizes our current information-saturated environment.*

— Expert

## **CONCLUSION: THE AI FLYWHEEL EFFECT**

Increased trust in AI systems and heightened dependence operate as mutually reinforcing mechanisms. As students develop greater confidence in the capabilities of AI, their usage broadens and deepens accordingly. The utility, efficiency, comfort, and ease of use of AI presents a particularly acute existential risk—when a technology can seemingly do everything, it becomes easier, even seemingly logical, for humans to do nothing (Burns 2024).

**BOX 13**

## Addressing students' AI dependence

Addressing students' dependence on AI requires building their awareness and agency while providing both alternatives to AI use and clear guidance about when and how to use it appropriately, as illustrated in the following examples:

*Awareness:* Teachers can build student awareness about AI overreliance through open discussions on dependence and its cognitive impacts while helping students to develop self-regulation skills and the forethought, focus on performance, and reflection associated with this (Zimmerman 2011). Students can complete parallel assignments with and without AI, then analyze differences in process, time, emotional state, and quality—making visible which cognitive capabilities they've surrendered to AI systems (Delphi panelist).

*Agency:* Through daily AI diaries, students can track when, why, how, and how much they use AI tools, then reflect weekly on its impact on their attention, confidence, and critical engagement. This consciousness-raising exercise can help students distinguish between outsourcing and augmenting their own capacities and can serve as an important step in reducing overuse of AI platforms (Delphi panelist).

*Abstinence:* Schools and families can embark on "digital detoxes" or "AI-free zones"—creating a temporal or spatial separation from AI use. Schools can provide AI "opt-out" assignments and activities to ensure that students maintain a healthy balance of AI and non-AI activities.

*Advisories:* AI platforms should display warnings when users exceed research-based thresholds, such as 15 inquiries within 24 hours. These "stopping cues"—similar to the viewing prompts of streaming services—alert students to risks to their cognitive ability, executive functioning, and well-being, and could automatically close platforms after extended use (Delphi panelist, Wood 2020).

This escalating dependence generates what can be characterized as a flywheel effect, transcending academic boundaries and creating spillover effects that extend into personal decisionmaking and daily life management (Figure 15). Parents and teachers confirm that this phenomenon has already materialized, with students reportedly using AI, not just for homework, but for "everything"—"relationships, entertainment, advice, life decisions, and mental health support."

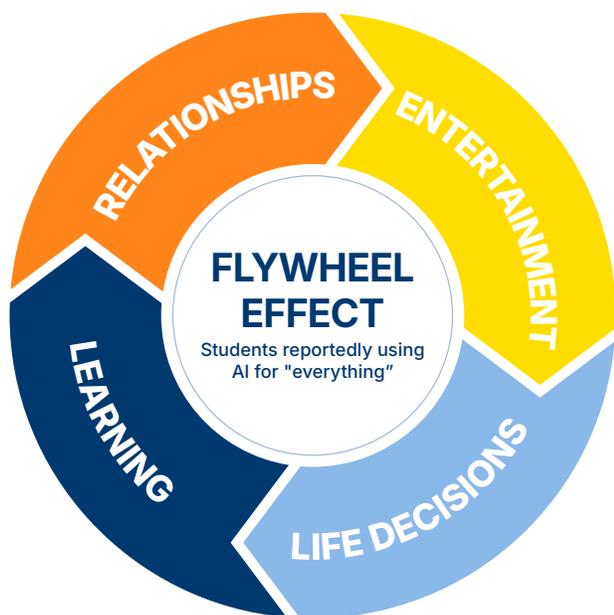
Data from the CDT confirms these observations. Surveys of student use of AI in the U.S. show an increasing range of student uses of AI from schoolwork (66%) to helping with relationships

(43%) to friendships (42%) and even romantic relationships (19%) (Laird et al. 2025). This expansion of AI from an academic tool to a comprehensive and essential life assistant holds serious implications, not just for students' education but for their personal development and growth.

As discussed, the dangers of overreliance for children and youth are manifold. Dependence may create feelings of inferiority, self-doubt, and a loss of confidence and self-efficacy—a belief that we cannot complete tasks or perform at a high level without use of AI tools—feelings that reinforce continued overuse (Burns 2024). It can result in a loss of agency and autonomy and potentially

FIGURE 15

## The AI flywheel effect on children's learning and development



© Center for Universal Education at the Brookings Institution, 2026.

a decrease in the risk-taking necessary for both learning and the formation of attachments and, as discussed in Risk 2, emotional attachments that can result in “heightened loneliness, reduced socialization, and increasing emotional dependence on the chatbots” (Fang et al. 2025). Dependence plays a factor in the degradation of trust that is foundational to teacher-student relationships.

But it is student learning that may bear the greatest brunt of excessive reliance on AI tools. As discussed in Risk 1, overreliance on AI can erode student judgment, autonomy, intellectual engagement, creative agency, and human capacity for ethical reflection and nuanced problem-solving (Bozkurt et al. 2024). For many students, this erosion has already happened. Furthermore, like any dependence, AI users may build up a tolerance to its profound risks and become habituated to its use (Burns 2024, 67). When users—in this case, students and their teachers—repeatedly delegate critical tasks to AI, they risk becoming passive

participants in decisions that demand moral judgement and emotional understanding (Bozkurt et al. 2024). It is at this point that agency becomes increasingly difficult to maintain, and it becomes harder to push back on this dependence on AI tools (Burns 2024).

Much of this overreliance can be mitigated with the proper institutional policies and practices and appropriate pedagogies, as will be discussed in Section VI. However, some of it cannot. The danger of highly effective and transformational tools lies in how their very designs make overreliance almost inevitable. Patterns of overreliance and dependence may ultimately lead to surrogacy—relinquishing control over complex outcomes that should remain in human hands and externalizing thinking and intellectual engagement to AI tools and systems, forsaking what Ralph Waldo Emerson called the “integrity” of the human mind (Burns 2024; Bozkurt et al. 2024; Emerson 1908, 15).

## Risk 6: AI can deepen equity divides

*Big Tech continues to develop solutions with the richest schools and students in mind.*

— Teacher

*The West is just extracting our [countries in sub-Saharan Africa] minerals and extracting our data.*

— Expert

The risks discussed above—from cognitive offloading to dependence on AI—primarily affect communities with access to AI. Early-adopter communities now grapple with overreliance on AI-generated content, student misuse for academic dishonesty, erosion of critical thinking skills, and ethical concerns around data privacy. Paradoxically, those traditionally excluded from technological advancements through lack of connectivity,

devices, or dominant language support may currently avoid these issues.

This adoption lag and access disparity may give late-adopter countries and communities valuable time to learn from early implementation challenges and develop more thoughtful policies, pedagogical frameworks, and support systems before widespread AI uptake. However, while such countries and education systems may avoid many of these current AI risks, they also miss out on potential benefits that early adopters are already experiencing. And, as we know from decades of technology provision in schools across the globe, many countries and education systems may remain entirely excluded from such benefits.

At present, AI's advantages accrue mainly to students in well-resourced countries and education systems, risking the exacerbation of existing inequalities. Like previous technologies, AI's benefits and risks are unevenly distributed across socioeconomic, gender, linguist, cultural, and geographic dimensions. Yet unlike earlier innovations, AI threatens to exacerbate these inequalities at unprecedented scale and in ways that are increasingly difficult to detect and address.

## THE AI DIVIDE

A critical finding emerging from the analysis of data to date is that a large and growing AI divide is widening and deepening what has long been an existing digital divide. This divide operates through a self-reinforcing cycle: as AI integration deepens within educational systems, already-existing barriers to accessing technology become increasingly entrenched. Beyond the perennial question of access, the AI divide introduces novel dimensions to existing disparities—algorithmic literacy, susceptibility to manipulation, and the capacity for critical evaluation of AI-generated content. This can lead to further sorting of winners and losers in which high-resource communities gain access to increasingly powerful AI technologies and deploy them productively, while low-resource communities either lack access

entirely or use these tools in less beneficial ways. This dynamic creates educational disparities that are both more pronounced and more difficult to remediate than previous technological divides.

This AI divide is evident across many dimensions but is especially visible by socioeconomic status (financial resources), geography (rural versus urban and different regions of the globe), and language (linguistic groups).

## THE SOCIOECONOMIC AI DIVIDE

Adopting, integrating, and using AI requires substantial financial investment for education systems. It requires:

- initial capital costs (hardware, software licenses, infrastructure setup, initial training)
- ongoing operational costs (electricity, Internet connectivity, cloud storage, technical support, maintenance)
- human resource costs (professional development, content adaptation, quality assurance time, pedagogical integration effort)
- unanticipated costs (the time in labor for teachers to check the accuracy of AI outputs for both their work and student products)

Home use of AI by students, while less resource intensive, still comes at a cost in terms of electricity and connectivity.

The calculus is simple—wealthy countries and communities can afford these costs. Poor ones cannot—and cost was repeatedly cited by study participants as a significant barrier to AI adoption.

Education systems have long confronted the substantial expenses associated with technology-based investments. Since most education systems allocate the majority of their budgets to staff salaries, technology costs represent a significant financial burden, particularly in poorer contexts (Our World in Data 2025). Lower-resource systems often lack discretionary funding for technology or AI tools, creating additional adoption barriers. These

financial constraints result in uneven technology adoption across schools and disparate access to technology among students, negatively and disproportionately affecting students who are poor, who live in rural areas, who are disabled, and who are female (Burns 2021).

Across the globe, many communities still lack fundamental infrastructure, including electricity, Internet connectivity, and mobile broadband access. Even where there may be access, poor schools often have older equipment which may function poorly and make them more vulnerable to attack. Pre-existing non-technical systemic shortfalls—including inadequate capital funding, lack of technology skills and specialized staff, and limited to no funding for licenses, maintenance and refreshing hardware—render access to technology, and by extension, quality AI, difficult or impossible for students in many educational contexts.

Many of the tools themselves demand even greater investment. Study participants report that AI-powered adaptive learning systems can cost tens of thousands of dollars per year. AI systems may need integration with existing school district software, such as student information systems and learning management systems, which incurs additional costs. Education systems may find they need completely new technology because existing equipment is too old or degraded to run AI programs effectively.

## THE GEOGRAPHIC AI DIVIDE

Access to AI-enabled education varies dramatically *within* countries, with rural communities facing severe disadvantages due to gaps in basic infrastructure. Disparities in access to electricity, Internet, and mobile broadband continue to reinforce the divide between rural and urban communities and their ability to benefit from AI. As of 2023, 84% of the global population without electricity resided in rural areas, with sub-Saharan Africa accounting for the majority without access (World Bank Group 2025).

Globally, 81% of residents of urban areas use the Internet, compared to 50% in rural areas. In sub-Saharan Africa, where both urban and rural use of the Internet is lower than global figures, this urban-rural usage gap stands at 57% versus 23%, respectively, as shown in Figure 16 (International Telecommunication Union (ITU) 2023b).

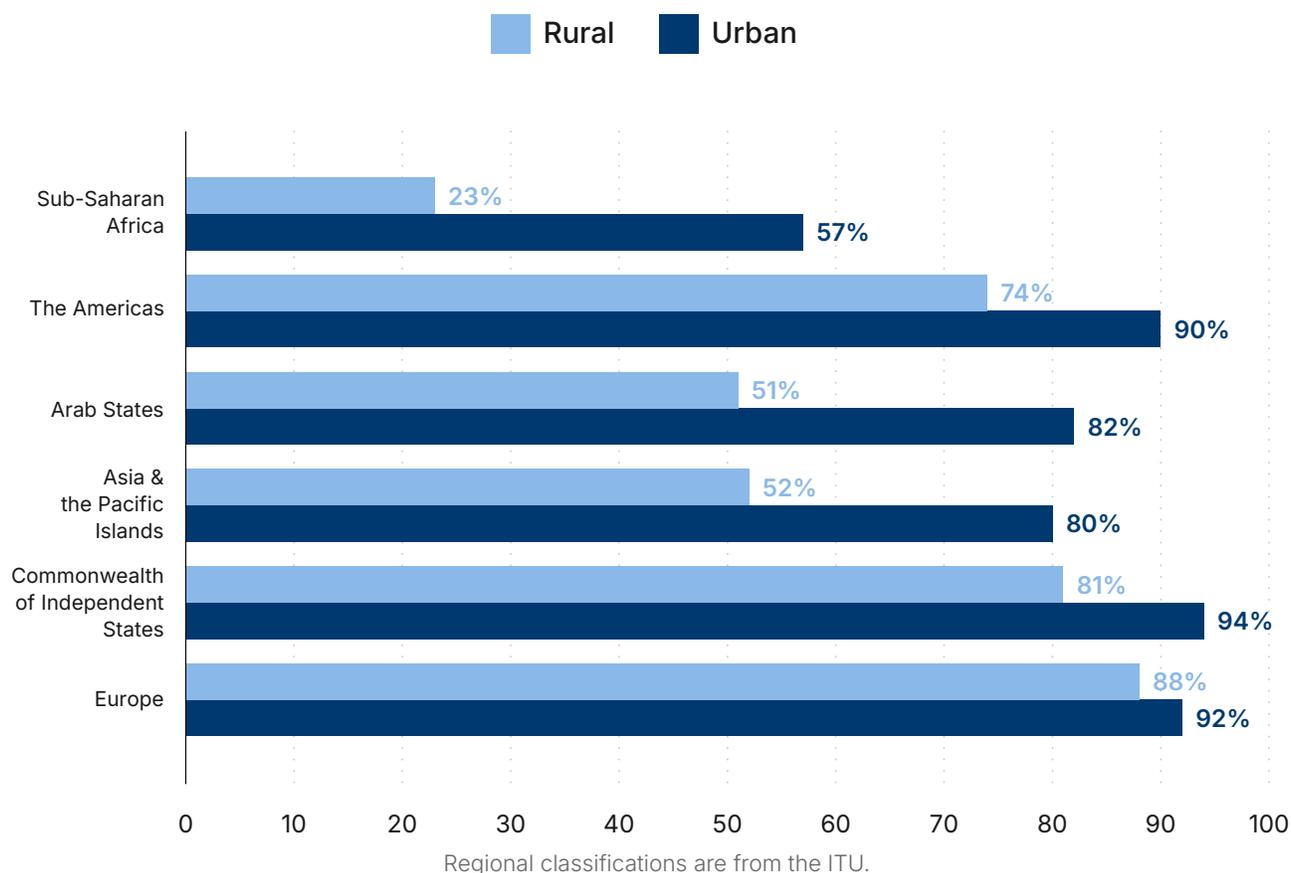
However, this rural-urban divide extends beyond access to AI tools and includes access to *information* and *policies* about AI use. In the U.S., Gallup and the Walton Family Foundation report that 67% of students in non-urban counties say their school has no clear AI policy; for students in urban areas, this figure is 50% (Gallup et al. 2025, 8). Clear AI use policies in school are positively correlated with increased confidence in using AI tools, meaning that rural students are at a disadvantage as they prepare for a workforce that increasingly expects AI skills (Gallup and Walton Family Foundation 2025)

The rural-urban divide is compounded by wealth since rural communities often have higher poverty rates and lower median incomes than urban ones (Su and Morgan 2024). The economic disparities within rural communities further limit AI access. In the U.S., students in lower-income rural counties are less than half as likely to be permitted to use AI at school compared to students in wealthier rural counties (15% vs. 41%). Gallup cautions that schools lacking clear AI policies and restricting AI use may be “hindering their students’ development of AI literacy” (Gallup et al. 2025, 9).

In addition to the rural-urban divide within countries, there is a broader AI divide across countries and regions. Many regions of the globe still lack fundamental infrastructure, including Internet connectivity and mobile phone ownership. For example, Internet usage stands at 66% in the Asia-Pacific region (comprising East Asia, Southeast Asia, Oceania, and parts of South Asia), 70% in the Middle East and North Africa, and 91% in Europe. Mobile phone ownership follows comparable patterns: 77% in Asia-Pacific, 83%

FIGURE 16

### Rural versus urban Internet usage



Source: ITU, 2023a, 2023b.

© Center for Universal Education at the Brookings Institution, 2026.

in the Middle East and North Africa, and 95% in Europe (ITU 2024).

Sub-Saharan Africa exemplifies these divides most starkly. Internet usage reaches only 38% in the region, while mobile phone ownership stands at 66% (ITU 2024). Most fundamentally, 600 million people—or 48% of the region’s total population—still lack reliable access to electricity. Thus, more than four out of five people worldwide lacking access to electricity live in sub-Saharan Africa (United Nations Sustainable Development Group 2025).

Compounding these access barriers, Internet costs in sub-Saharan Africa are among the world’s most

expensive. As of 2025, Ghana is the second most expensive country with the highest fixed broadband costs globally, with rates reaching approximately \$2.58 per megabit per second. Among the 10 countries with the most expensive fixed broadband Internet, three—Ghana, Kenya, and Nigeria—are located in sub-Saharan Africa. This stands in stark contrast to the least expensive markets, where countries such as Romania, Thailand, and Singapore offer fixed broadband Internet at costs as low as \$0.01 to \$0.02 per megabit per second (Venditti 2025).

Finally, in terms of AI specifically, the paucity of data centers in sub-Saharan Africa (less than 1% of all the world’s data centers) means that teachers

## BOX 14

# Data centers

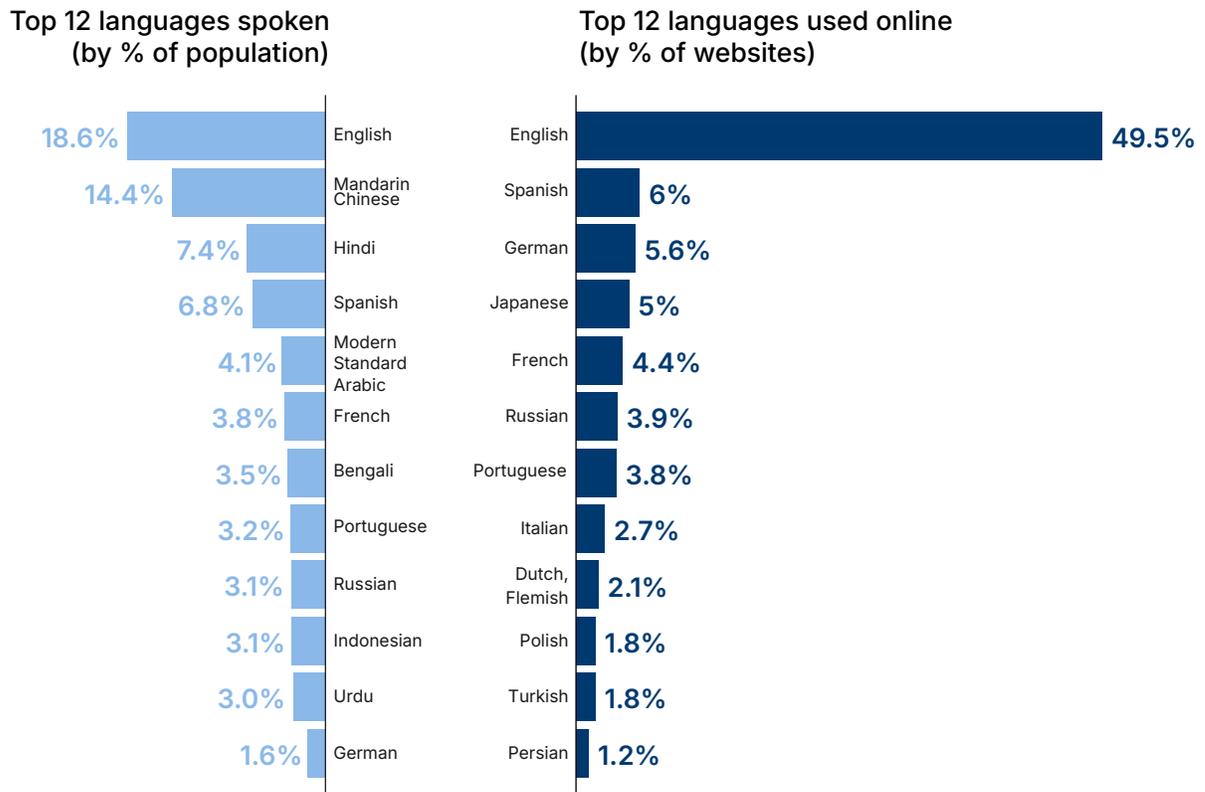
Addressing many of the layers of the AI divide—for example, access to AI platforms, tools, and content in local languages—requires data centers—which tend to be concentrated in wealthy countries. Data centers are specialized facilities that manage Internet technology infrastructure such as servers, storage devices, and network equipment. They play a critical role in processing, storing, and distributing large amounts of data, making them essential to AI and to the rest of the digital economy. But they also require substantial “compute” power—the hardware, processors, memory, storage, and energy needed to operate data centers effectively as well as electricity. At the time of writing, South Africa is the only African country where Amazon, Microsoft, and Google have built their own data centers (Dosunmu 2025).

and students must depend on foreign-developed AI models that are not designed to reflect local contexts, languages, curricula, or pedagogical approaches (Stasenکو et al. 2025; Bracewell 2025). While grassroots initiatives like Masakhane and Lelapa have successfully created African language models, the vast majority of AI training for African languages still occurs on foreign infrastructure (Elongué 2025). The consequences of this extend beyond linguistic issues to encompass data sovereignty concerns and limitations on the scale of locally developed solutions. Further, as UNESCO argues, they may represent a new form of “digital colonialism” manifested as “data poverty,” where communities lacking sufficient data find themselves “excluded and put at long-term risk of being colonized by the standards embedded in the GPT models” (Miao and Holmes 2023, 14). African technology specialists and policymakers interviewed for this study warn of a “vicious cycle” whereby the continued advancement of AI exacerbates the digital divide, leaving African students further behind and preventing the continent from developing the indigenous AI capabilities needed for culturally responsive educational technologies.

## THE LINGUISTIC AND CULTURAL AI DIVIDE

Language represents the variety of human cultural contexts, ways of seeing the world, and the experiences, values, and understandings of their speakers. There are over 7,000 living languages across the globe (UNESCO 2025). Yet one predominates in AI—English. Almost half of Internet content is in English (Statista 2025b). The dominance of AI tools in English disadvantages underrepresented communities (Okolo and Tano 2024). Of the top 20 languages used on the Internet, not one comes from Africa’s more than 2000 languages, as shown in Figure 17 (Statista 2025b; Ethnologue 2025; United Nations Population Division 2025). Hausa is spoken by 94 million people, primarily in Nigeria (sub-Saharan Africa’s largest country and the world’s sixth-most populous nation), yet ChatGPT recognizes only 10–20% of sentences written in Hausa (Moorosi 2024).

Top 12 spoken versus online languages



Source: Statista, 2025; Ethnologue <sup>5</sup>, 2025; United Nations Population Division, 2025.  
 © Center for Universal Education at the Brookings Institution, 2026.

This linguistic and cultural divide, in particular, is a result of the narrow development base from which AI tools and platforms emerge. Technology traditionally, and AI most recently, embodies the “WEIRD problem”—the overrepresentation of Western, English-speaking, Industrialized, Rich, and Democratic contexts in their design and training (Henrich et al. 2010). This bias means that AI disproportionately represents the data, belief systems, values, and languages of WEIRD societies while underrepresenting those of the majority of the global population. Educational technology developers prioritize affluent WEIRD markets over broad educational access and rely on training data that inadequately represents diverse student populations.

**THE AI ACCURACY DIVIDE**

The concentration of AI development among large technology companies and wealthy education systems deepens existing digital disparities in unprecedented ways and fuels a growing AI gap (Lanz 2025; Winsome Marketing Writing Team 2025). This emerging gap manifests as what one study participant called “cognitive stratification,” wherein accuracy of information and advanced reasoning become “premium commodities.” Wealthy educational systems can access sophisticated AI tools that provide more factual information, while the poorest schools, teachers, and students rely on LLMs that produce less accurate information.

Like most technology disparities, this cognitive stratification is rooted in cost and affordability. Discussions with educators in our study suggest that many schools, particularly in low-resource environments, rely on the free version of LLMs, overwhelmingly ChatGPT. Free tiers come with constraints such as usage limits, reduced features, slower response times, or access to older model versions. Critically, free versions are also much less accurate in the information they produce and in their reasoning capabilities (Lelièvre et al. 2025).

### BOX 15

## Measuring the accuracy of LLMs

AI accuracy is measured by benchmarks such as Massive Multitask Language Understanding (MMLUs), a standardized test for AI models that measures their knowledge across 57 academic subjects using multiple-choice questions. AI benchmarking is similar to standardized testing in education. It provides an objective way to evaluate and compare how well different AI models perform on specific tasks. Just as schools use test scores to measure student achievement, MMLU scores help determine which AI models are most accurate and reliable for educational use (Hendrycks et al. 2021).

LLMs are priced in tiers, including free and reduced-cost versions for schools. For example, at the time of writing, students and educators can access the free version of ChatGPT, an individual subscription of ChatGPT (\$20 per month), or a “team plan” which cost \$25–30 per user monthly (OpenAI 2025e). Some teachers interviewed reported that they pool funds to buy an upgraded license that they share. “Enterprise” versions can cost hundreds of dollars per user per month. Paid subscriptions typically provide access to models

with larger parameter counts, more extensive training, and superior reasoning capabilities.

For a secondary school of 1,000 students, full ChatGPT Plus access could cost \$200,000–300,000 annually. More specialized educational AI tools demand even greater investment, with adaptive learning systems costing tens of thousands of dollars yearly, plus additional expenses for integration with existing school software and potential hardware upgrades.

Importantly, in contrast to free versions of models, tiered systems cost more, but they do more and they are more factually accurate. As OpenAI itself has acknowledged, all models of its ChatGPT platform hallucinate—that is, produce false or nonsensical information—regardless of cost (OpenAI 2025g). However, tiered models tend to offer greater functionality and perform more accurately across a broader range of tasks, including reasoning, factual accuracy, and nuanced understanding, particularly on standardized benchmarks (OpenAI et al. 2024). The functionality (and, in most cases, the accuracy rates) improve commensurate with pricing (Hendrycks et al. 2021).

Accuracy of information produced depends on the domain and tasks, so accuracy rates must be considered as illustrative versus absolute. In addition, the relationship between cost and accuracy is not absolutely linear, nor is it as simple as “free” versus “paid.” However, the relationship between cost and accuracy is largely consistent. Research comparing model performance across pricing tiers indicates that free models generally demonstrate lower accuracy than paid alternatives, though this does not necessarily mean that the most expensive ones are also the most accurate (Lelièvre et al. 2025; Hardman 2025). This performance differential is most pronounced in complex reasoning tasks and factual accuracy measures, where paid-tiered models, like GPT-4, demonstrate superior performance across professional and academic benchmarks over their free-tier counterparts such as GPT-3.5 (OpenAI et al. 2023; Hardman 2025). For example, ChatGPT

---

## This may be the first time in the history of educational technology that schools must pay more for accurate factual information.

---

3.5, which is still in use in many parts of the globe at the time of writing, has an accuracy rate of 70% on most MMLU benchmarks, while tiered versions report 86% accuracy rates (OpenAI et al. 2023).

But many schools around the world can't afford to pay for even discounted use of LLMs. Thus, relying solely on free tiers, as many of these schools, students and teachers do, compounds existing educational disparities. Schools have long paid more money for technology tools with greater functionality. However, this may be the first time in the history of educational technology that schools must pay more for *accurate factual information*.

### The Matthew Effect

Taken together, these divides—socioeconomic, rural-urban, regional, linguistic and educational—accentuate a longstanding pattern in educational technology known as the Matthew Effect, where the “rich get richer,” with wealthy students who use AI effectively capitalizing on its benefits while poor students do not (Mishra 2024).

The AI divide splinters into more granular dimensions that affect how technology functions within educational contexts. A *capacity divide* reflects uneven distribution of abilities to use AI effectively, encompassing both strategic capacity (using AI to achieve specific goals) and productive capacity (the formal skills needed to engage with AI systems) (van Deursen and van Dijk 2014). A *usage divide* emerges as poorer schools deploy technology to reinforce lower-order, rote-based learning rather than higher-order thinking, while a *benefits divide* means that even when poor students access technology, they reap fewer benefits compared to their affluent peers

(Burns 2021, 59; van Deursen and van Dijk 2014). A *vocational divide* further compounds these inequities as increasing percentages of future jobs will require exceptional AI skills while others will be eliminated.

Educational technology has long perpetuated this Matthew Effect, but AI risks widening the gap dramatically. Students with strong foundational knowledge leverage AI to deepen learning and progress faster, while struggling learners, lacking expertise, may misinterpret AI outputs and fall further behind (Mishra 2024; Kaufman et al. 2024). These divides exacerbate pre-existing structural inequalities, threatening to transform AI's promise of democratizing education into a mechanism that entrenches existing privilege.

### CONCLUSION: IT'S NOT TOO LATE TO ADDRESS THE AI DIVIDE

These equity challenges should not push us into a cycle of despair. Internet use has increased over the years, including in sub-Saharan Africa, thanks to more affordable plans, better infrastructure, and concerted multilateral efforts to expand access, such as the World Bank Group's Digital Economy for Africa (DE4A) initiative (World Bank 2023; World Bank 2024).

Equity is often more of an issue of political will versus an engineering problem, and there is much that governments, multilateral aid agencies, and technology companies can do to bridge the AI disparities discussed here. As Section VI will discuss, governments can undertake any number of innovative financing and regulatory schemes to ensure greater technology access.

**BOX 16**

## From policy to implementation: Addressing the AI divide

Realizing AI's transformative potential for educational advancement across all nations requires developing local AI ecosystems. This in turn requires coordinated action spanning policy frameworks, sustained investment, and operational implementation—a multidimensional challenge that diverse stakeholders are increasingly addressing through complementary approaches:

At the *policy level*, the Hamburg Declaration on Responsible AI for the Sustainable Development Goals represents the first global declaration focused specifically on AI in international development, outlining shared principles and commitments to promote equitable, inclusive, and sustainable development and deployment of AI worldwide, with particular emphasis on empowering developing countries. The declaration establishes equity as its foundation, characterizing AI as a “digital public good” and calling for expanded educational access for women and girls alongside support for local AI innovations in developing countries (United Nations Development Programme 2025).

But translating these principles into practice requires substantial *financial commitment*. The Gates Foundation's investments across Africa exemplify such support, including \$30 million to develop safe, ethical, and equitable AI platforms designed to “pave the way for local innovation” and \$7.5 million to Nigeria's Artificial Intelligence Scaling Hub to “promote the ethical development and widespread adoption of AI technologies in sectors, including healthcare, agriculture, and education” (Empower Africa 2025; Zamora 2023). These initiatives represent the financial support that Global South governments need to create their own AI ecosystems.

Beyond policy and funding, *sustained implementation mechanisms* are essential for an AI ecosystem to develop and thrive. The United Nations Development Programme's AI Hub for Sustainable Development in Africa has accelerated innovation by connecting external funders with local AI talent in sub-Saharan Africa and improving access to representative datasets, increasing computing power availability, attracting and retaining technical talent, prioritizing local contexts, and addressing regulatory frameworks (Hradecky and Abdulkareem 2025).



## V **RISKS VERSUS BENEFITS:** Currently, AI's risks overshadow its benefits

Will AI transform education to help all children thrive—or generate harms that undercut potential gains?

The answer, as we've seen, is that it may do both. Ultimately, it depends on how AI is integrated into the heart of teaching and learning.

### **AI's benefits are numerous**

This report has highlighted numerous examples of how AI enriches and extends learning when used as a supplement rather than a substitute for human instruction. AI can free teachers from administrative tasks, enabling more high-value interactions with students. Personalized learning and individualized tutoring can deliver effective one-on-one instruction, improving academic performance at scale through immediate feedback, adaptive pacing, and customized explanations.

AWE systems help students produce coherent, well-structured essays, with linguistically disadvantaged learners showing the greatest gains—suggesting AI can narrow language gaps and support second-language acquisition.

AI can also adapt learning and assessment to student ability, bridging gaps between students from different educational backgrounds, particularly where teacher resources are limited. Social chatbots offer neurodivergent learners safe opportunities to rehearse social interactions without fear of judgment, while assistive technologies expand access for students with disabilities. When combined with predictive AI, generative AI creates powerful synergies: predictive systems identify learning needs and generative tools respond with personalized interventions.

## **But today, these benefits are overshadowed by risks**

While these benefits are significant, both existing research and participant data suggest that, under current conditions, AI's harms overshadow its advantages. These risks differ in nature from the benefits, undermining the very learning capacities and relationships needed to leverage AI's strengths.

Without safety guardrails and bounded learning content, AI can encourage cognitive offloading—students delegating complex thinking to machines. This reliance erodes critical thinking, weakens content knowledge, blurs fact from falsehood, and diminishes communication and durable skills essential for thriving in an AI-infused world.

Emotionally, AI companions exploit vulnerabilities through anthropomorphism and simulated empathy, fostering unhealthy digital attachments—especially among vulnerable students. The servile design

of chatbots, their sycophantic tone, vast data-processing capabilities, and frictionless interactions can hinder social skill development and enable manipulation, hyperpersuasion, and polarization. Deepfakes, academic dishonesty, concealed AI use, acceptance of hallucinations as truth, and the “black box” nature of AI further erode trust in schools and institutions.

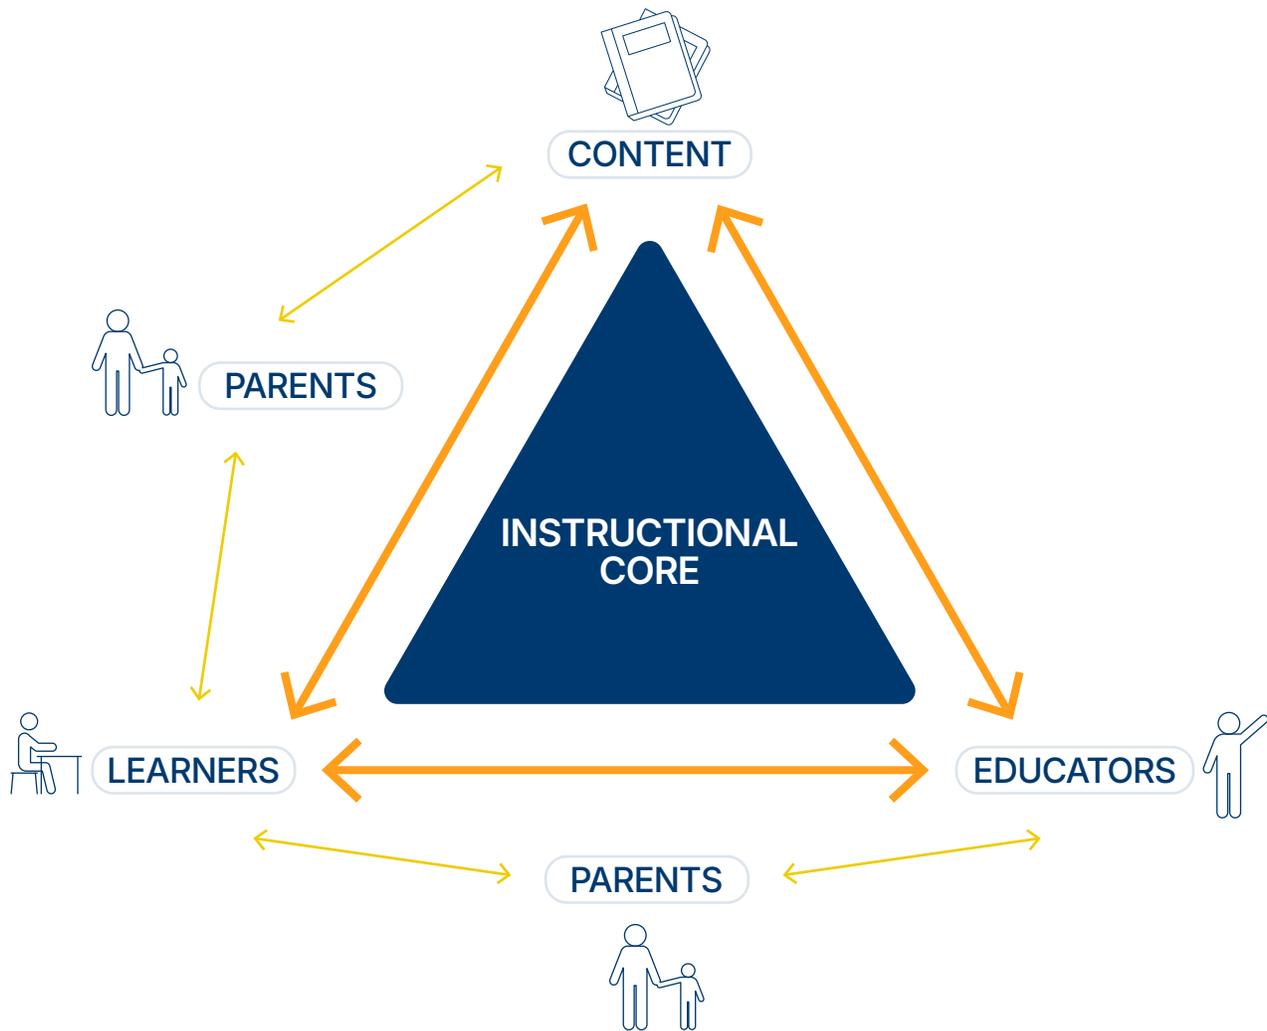
AI tools also pose serious privacy and safety risks, including data exploitation, exposure to explicit content, nonconsensual deepfake imagery, and reinforcement of bias and polarization. Moreover, AI amplifies advantages for students with strong foundational knowledge while widening achievement gaps for underprivileged learners.

## **Advancing AI learning experiences that enrich, not diminish**

The harms identified in this report are neither inevitable nor immutable. The key question is not what AI can do, but what it should do within the context of supporting children's cognitive, social, and emotional development.

AI implementation must be designed to enrich—not diminish—learning. Decades of research show that the interaction between students, teachers, and content forms the instructional core that drives educational change (see Figure 18) (Cohen and Ball 1999). Including parents in this core acknowledges that learning ecosystems extend into homes and communities, meaning AI design should consider its impact on parent-child interactions, parental insight into student progress, and family engagement with educational content (Vegas and Winthrop 2020).

The instructional core that drives educational change



Source: Vegas and Winthrop, 2020.

© Center for Universal Education at the Brookings Institution, 2026.

When AI strengthens the instructional core—by expanding the capabilities of key actors or improving their interactions—it enriches learning. The reverse is also true: when AI erodes these capabilities or relationships, it diminishes learning.

Many benefits discussed above and in Section III enhance the instructional core. For example, AI can create synthetic voices for students with aphasia, unlocking new ways to communicate with teachers. Embedded prompts in digital materials can help second-language learners through interactive

question and answer functions, opening new pathways for engagement. AI-powered assessment tools can give teachers nuanced insights into student learning, enabling more targeted instruction.

Conversely, the risks outlined above and in Section IV threaten these same capabilities and relationships. AI diminishes learning when it undermines trust among students, teachers, and parents; when it exposes students to bias, manipulation, and privacy violations; and when it

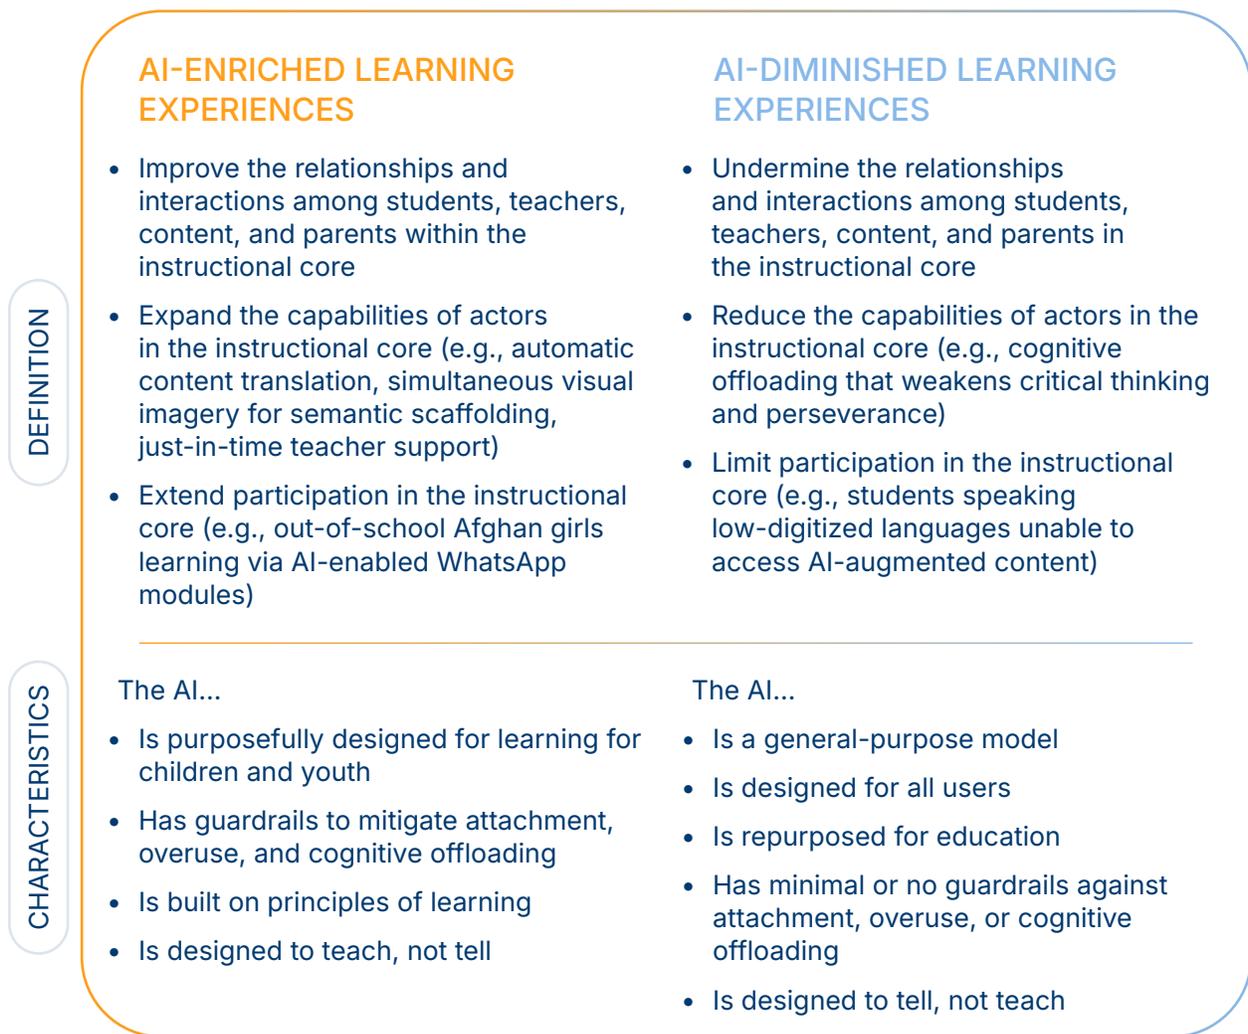
accelerates disengagement or excludes certain groups from AI's benefits.

Our findings suggest that AI especially enriches learning when its use is bounded by vetted content and safety guardrails within structured learning activities. In contrast, it tends to diminish learning

when students engage with general-purpose AI tools in open-ended, unregulated contexts. Figure 19 summarizes these patterns, offering a high-level view of when and how AI use typically enriches—or undermines—learning. These use cases will evolve as AI develops, but they provide a starting point for clarifying today's debates on minimizing risks while harnessing benefits.

FIGURE 19

### AI enriched versus diminished learning experiences



© Center for Universal Education at the Brookings Institution, 2026.

## Drivers of AI risks

Three key drivers shape how AI impacts children and help explain why current AI implementations may diminish learning experiences.

### PSYCHOLOGICAL AND EMOTIONAL DEVELOPMENT

Children are uniquely vulnerable because their brains, knowledge, and emotions are still forming. Human tendencies toward anthropomorphism make them highly susceptible to AI's simulated social cues—an effect amplified by modern technology (Epley 2019). Convenience, shortcuts, and grade-driven feedback loops encourage cognitive offloading and dependency. AI companions exploit emotional vulnerabilities through emulated empathy and mirroring, which can hinder social skill development and deepen reliance on artificial relationships.

### AI TOOL DESIGN

AI does not possess true intelligence; it operates through statistical pattern recognition rather than reasoning or comprehension. Consumer tools prioritize speed and engagement over safety or learning, emphasizing automation instead of critical thinking. They frequently hallucinate—confidently presenting false or nonsensical information. Student-facing tools often provide an “illusion of impact”: assumed to be high-value but frequently modeling poor pedagogy, misunderstanding how children learn, and perpetuating rote approaches (Hardman 2025; Chen et al. 2025).

### PEDAGOGY

Research shows overreliance on AI harms learning, with dosage—the extent and frequency of use—emerging as critical. Benefits occur when teachers integrate AI thoughtfully to enrich, not replace, traditional methods (Trabelsi 2025). Yet instructional design with AI is often absent; many educators simply layer AI onto existing

frameworks without reimagining teaching around these tools. Without deliberate redesigning, AI risks undermining core educational goals rather than enhancing them.

## The time to act is now

AI promises both unprecedented opportunity and significant threats. There is much to gain from generative AI tools that enrich education, but there is also much to lose when AI tools erode what is foundational to teaching and learning. All study participants—particularly parents and teachers—speak with one voice: the time to act is now. We turn next to recommendations for action to minimize the risks and maximize the benefits of AI for all children and youth



## VI IT'S NOT TOO LATE: Recommendations for mitigating risks and harnessing benefits

As this report had made clear, AI is reshaping education for students across the globe, with outcomes that remain uncertain. Without proactive and comprehensive intervention, the risks and harms we have identified are likely to persist and intensify, potentially outweighing the benefits of AI and widening disparities in access and effective use. But it is not too late to shift the trajectory of AI in education away from uses that diminish learning and toward those that enrich it.

We must act with urgency. Governments face a narrowing window to develop protective regulations, while educational institutions must quickly establish safeguards and pedagogical approaches that maximize benefits and mitigate harms. Immediate protective measures must include protocols for identifying sycophantic AI responses, recognizing misinformation and bias, and ensuring sustained human oversight of AI-mediated learning. Curricula must teach students to critically evaluate AI-generated content, understand AI's limitations, and preserve their own intellectual agency. Those working with children need

---

## The actions we take—or fail to take—now will determine whether AI becomes education’s greatest asset for student learning and wellbeing or its most profound threat to student flourishing.

---

holistic AI literacy education—covering how AI functions, its cognitive and emotional impacts on young people, strategies to reduce cognitive harms, and techniques for verifying information.

These multifaceted safety and security concerns require good-faith participation from everyone: legislators, policymakers, AI developers, educators, civil society organizations, families, and students. Most importantly, AI companies, governments, and education systems must collaborate to ensure the responsible design of education-facing AI tools, enabling the full power of AI’s benefits to be harnessed for all students and schools. The actions we take—or fail to take—now will determine whether AI becomes education’s greatest asset for student learning and wellbeing or its most profound threat to student flourishing.

The recommendations that follow, drawn from study participants, Delphi panelists, and a growing body of research on AI in children’s education, outline practical and preventive steps necessary to protect students from AI-related harms while harnessing AI’s benefits and preparing young people to thrive in an AI-suffused world. These measures aim to ensure that all students can grow through educational experiences that cultivate learning, agency, and future success.

### **PROSPER, PREPARE, PROTECT**

To seize this opportunity, we all have a role to play. Our proposed framework rests on three pillars—Prosper, Prepare, and Protect. Each pillar is discrete yet equally vital, but all three must work in concert to ensure that AI supports children’s flourishing, safeguards their well-being, and equips them to meet the challenges of tomorrow.

## A framework for action



**PROSPER:** Recommendations under the Prosper pillar focus on transforming teaching and learning experiences so that children and youth can thrive in an education system where AI is omnipresent. Students can *prosper* through carefully titrated AI use (knowing when to teach with and without AI, using AI only when it enhances rather than replaces student effort and cognitive engagement), high-quality pedagogical integration (combining AI with evidence-based practices that prioritize deeper learning), and collaborative design and research (co-designing tools with educators and communities while conducting rigorous research on when and how AI supports learning). These efforts will help children and youth develop the capabilities they need to flourish in an AI-driven world.

**PREPARE:** Recommendations under the Prepare pillar focus on building the knowledge, capacity, and structures needed for students, educators, families, and education systems to integrate AI ethically, effectively, and humanely. The education community can *prepare* for a rapidly evolving AI-infused world through holistic AI literacy (developing understanding about AI's capabilities, limitations, and implications), robust professional development (equipping educators with knowledge and skills to teach with and about AI), and systemic planning and access (establishing clear visions for ethical AI use while expanding equitable access). These efforts will ensure that students, educators, families, and education systems can navigate the world of AI with confidence and responsibility.

**PROTECT:** Recommendations under the Protect pillar include developing and implementing safeguards on AI for student privacy, safety, emotional well-being, and cognitive and social development. Students can be *protected* through ethical and trustworthy AI design (protections embedded into the technology during the design phase), responsible governance (strong regulatory frameworks), and adult guidance (modeling healthy technology use at home and in schools). These efforts will help ensure that AI enhances learning while keeping children safe, supported, and able to thrive in an AI-infused world.

# **We all have a role to play**

To help mitigate risks and harness benefits, all actors touching AI and education will need to be involved. The following stakeholders play a critical role in designing, implementing, enacting and monitoring our recommendations.

---

## **GOVERNMENTS**

- » Government actors include federal/national, state/provincial, local and tribal/indigenous governments. The scope and levers of these government actors vary across the globe. Local context and conditions will inform decisions about which jurisdictions are best able to implement these recommendations.

---

## **EDUCATION FUNDERS**

- » This includes foundations, private sector donors, corporate social responsibility actors, philanthropists, and government funders who provide financial backing to education programming.

---

## **TECHNOLOGY COMPANIES**

- » This includes general-purpose AI, social media, AI in education and educational technology companies who develop, adapt, market, or sell AI products to education systems or directly to children and youth under the age of 18.

---

## **EDUCATION SYSTEMS**

- » Education systems encompass multiple institutional levels, from the macro to microlevel. They include institutions such as ministries of education, state and regional education offices, schools, and universities. These institutions are composed of stakeholders who have decisionmaking power within these systems. They may include superintendents, regional and district supervisors, teachers' unions, school leaders, head teachers, teachers, and paraprofessionals. The structure of education systems varies from country to country, therefore local contexts should be considered to determine how best to enact each of the recommendations we outline below.

---

## **CIVIL SOCIETY ORGANIZATIONS**

- » These include nongovernmental organizations, faith-based organizations, and others that work to advance children's education and shift policy and practice.

---

## **FAMILIES, CAREGIVERS, AND COMMUNITY ACTORS**

- » This includes all adults that interact with or impact children outside of school—in the home or broader community. This may include parents, caregivers, health care providers, childcare providers, and social workers, among others.

---

## **RESEARCHERS AND ACADEMICS**

- » This involves anyone conducting research related to education in an AI world, such as researchers at universities and other institutions of higher education, think tanks, civil society, funders, the private sector, education organizations, and program implementers.

---

## **STUDENTS**

- » Finally, and most critically, children must be active participants in shaping how AI is used responsibly and ethically in educational settings. They must be allowed to contribute their perspectives and

experiences to inform the policies and practices that affect their learning and must understand their own role in responsible AI use.

Across our three pillars, we present 12 concrete actions that, if enacted, will minimize the risk of AI diminishing learning and maximize its potential to enrich it. Some of these actions can be undertaken by a broad range of actors, while others require the involvement of specific stakeholders. For each action, we list all the relevant actors who can help advance it.

We encourage readers to choose at least one recommendation—and preferably several—to which they can dedicate time, resources, and attention, and then take concrete steps to advance them.

## | Prosper

Prospering involves transforming teaching and learning experiences so that children and youth can thrive as learners, professionals, citizens, and human beings in a world where AI increasingly permeates every dimension of their technological, vocational, civic, and personal lives. To ensure that learners prosper, governments, funders, the private sector, education systems, families, and communities can work together to advance these four main actions:

1. **Shift educational experiences in school.** Perhaps one of the most promising opportunities AI presents is transforming teaching and learning experiences in school to align with the approaches that educators and learning scientists have long championed. This will require intentionally identifying where, when, and how AI should and should not be used in education.
2. **Co-create educational AI tools with educators, students, parents, and communities.** One of the best ways to ensure that AI tools are fit for purpose and advance education's objectives is to co-design them with users, including those in marginalized communities.
3. **Use AI tools that teach, not tell.** Although ed-tech companies may incorporate the principles of child-focused learning sciences when developing educational tools, general purpose technologies—because of their universal audience—typically do not. Yet general purpose tools like LLMs are the dominant technology in our study schools. Thus, AI companies with products students access should work to ensure that these AI tools are better optimized for children's learning needs and developmental stages.
4. **Conduct research on children's learning and development in an AI world.** Education policymakers, program designers, and educators need more data on AI's role and long-term impact on student learning, particularly on students' cognitive skills, affective states, and motivational states. Salient questions that researchers can investigate are: What are AI's effects on learning? Are these impacts the result of the inherent qualities of AI tools or of how they are used (Deng et al. 2025, 39)? What specific AI tools and use cases best help students learn? When should AI be used, and should it be avoided? How can AI enrich education, extending and deepening the teaching and learning process for all students?

## 1 SHIFT EDUCATIONAL EXPERIENCES IN SCHOOLS

Perhaps one of the most important steps enabling students to prosper in an AI world is to shift education experiences away from the transactional task completion that characterizes learning in so many schools worldwide. There is no shortage of ideas on how to shift what David Tyack and William Tobin call “the grammar of schooling,” namely the central structure of most education systems (Tyack and Tobin 1994). Schooling models that focus on meaningful practice and deeper learning have been advanced for decades, even centuries, worldwide. These range from community-based learning circles to vocational exploration that integrates academic learning (Dewey 1916; Freire 1970; Krishnamurti 1953). Today, many of these approaches have the added benefit of being supported by emerging learning science. For example, students learn better when they feel they belong, when they get a chance to practice and apply what they learn, and when they are focused and emotionally invested in what they are doing in school (Cohen 2022; Darling-Hammond et al. 2020; Immordino-Yang and Damasio 2007).

Many of the harms identified here—particularly the risks to students’ learning—originate largely from attempting to overlay transformative technology onto educational structures that have, at their core, remained largely unchanged since the late nineteenth century (Tyack and Tobin 1994). As one secondary school student described it, the core student schooling experience is “memorize, recite, forget,” year after year. While AI should not drive educational change, it lays bare weaknesses in current systems and provides education systems with a strong motivation to reform their purposes and processes.

Schools and extracurricular and out-of-school programs can coordinate efforts to help young people retain and develop their humanity as individuals and as members of their communities. This includes identifying when and how AI should, and should not, be integrated into teaching and learning experiences. *Education systems, governments, families, students, civil society organizations, technology companies, funders, and researchers* all have a role to play in helping young people have learning experiences that help them develop the knowledge, competencies, and skills they need to thrive in an AI world.

National and multi-country initiatives that help *education systems* learn from each other on how to shift teaching and learning practices can play an essential role in accelerating these shifts. Existing initiatives can be funded to deepen or expand their work, including focusing on the four areas of learning below. In some cases, new initiatives will need to be established where existing efforts fall short. To build momentum, insight, and sustainability, these national or multi-country initiatives can involve an array of actors involved in supporting education alongside education systems.

To shift educational experiences in schools, **craft new pedagogies that are AI-aware, AI-assisted, and, when necessary, AI-resistant.**

Technology is not pedagogy. The introduction of AI into classrooms, no matter how powerful, does not in itself constitute instruction or innovative pedagogy. Critical cognitive skills like deep reading, critical thinking, and Level 3 and 4 writing are hard to teach and resist simple technological solutions. While AI-assisted pedagogies can deepen students’ engagement with these areas, AI cannot support these skills unless the instructional practices around them are improved first. This includes understanding effective reading-writing connections, building classroom environments that support collaborative writing, addressing the needs of diverse learners, and resisting reductive practices tied to high-stakes testing.

Research continues to show that motivation, formative feedback, and access to adequate resources are all crucial, and these cannot be outsourced to machines (Graham 2019).

Teachers can identify what aspects of teaching and learning are most vulnerable to automation or distortion by AI. For example, they can distinguish between tasks where AI can scaffold learning—such as drafting, idea generation, or revising writing—and those that demand human interaction, ethical judgment, and the cultivation of learner identity. These insights should guide AI-assisted decisions (e.g., leveraging AI for feedback or modeling writing techniques) and AI-resistant choices (prioritizing human dialogue, peer collaboration, and cultivating students' own voice). This strategic approach leverages AI to unlock new educational possibilities rather than simply accelerating existing practices (Delphi panelist).

Ultimately, teachers can craft pedagogies that are both robust in the face of technological change and expansive in their vision of human learning. AI should be an ally in this work—not as the driver of learning, but as a tool in the hands of skilled teachers who understand how learning truly happens (Delphi panelist).

| To shift educational experiences in schools, **develop students' core learning capacities.**

Around the world, children need experiences in and out of school that support the development of a range of foundational skills essential for learning. These skills, often described as a suite of brain processes, include things such as self-control, working memory, selective attention, and cognitive flexibility (Diamond 2013). Children develop these skills starting from the moment they are born up until and through early adulthood. These skills can be taught, improve with practice, and are essential for children's cognitive, social, and emotional development. They help children succeed in school and life, from learning to read to addressing unexpected challenges, and from seeing other people's perspectives to solving problems (Diamond 2013; Galinsky 2024).

The potential risk, that poor AI use may undermine these foundational learning capacities, can be mitigated through shifting teaching and learning experiences, including employing the actions in the Box 17 below.

## BOX 17

# Strategies to support students' core learning capacities

**Playful and interactive learning activities.** Educators can help students develop core learning capacities by incorporating playful and interactive pedagogical approaches into their regular lessons. For the early years, this can include strategies such as students practicing reading to each other, with one student listening (and not interrupting), then asking a question after the reader is done. This helps children develop their self-control, attention, and working memory. Collaborative storytelling, where children play together, including taking on roles in a pretend scenario (such as playing grandma cooking at home), helps develop cognitive flexibility. From the Tools of the Mind intervention in the U.S., to group play and physical activity in China, to cards and table top games in Spain, there are a wide range of ways educators can incorporate these activities into their regular lessons (Diamond and Lee 2011; Bodrova and Leong 2007; Bai et al. 2022; Vita-Barrull et al. 2023). As children grow, these skills

develop further. Once-in-time tests that ask students to recall information are much less helpful in developing these essential capabilities than when educators ask students to engage in ethical dialogue, debates (including switching sides midway through), or when students undertake tasks that require them to plan, monitor, and adapt their work (Alexander 2018; Immordino-Yang 2025).

**School phone bans with a pedagogical exception.** Cell phones provide a host of distractions to young people and can interfere with children’s attention. With the integration of AI into virtually all technology platforms, including social media, the ongoing call to limit student cell phone use in schools will continue to be important. Many school systems have already instituted limited student cell phone use, including “bell-to-bell” phone bans that ensure young people spend time interacting in person during mealtimes and breaks (Economist 2025). However, for low-resource schools, including those without science labs or access to other technology, phones may be the one access point where educators can provide useful technology-based learning opportunities, including learning about AI. In these contexts, limiting student phone use throughout the day, except for legitimate pedagogical use, when sanctioned by educators, is helpful for students social as well as academic development.

To shift educational experiences in schools, **develop students’ knowledge and understanding, through interdisciplinary frameworks.**

Education systems must recognize the humanities and social sciences as essential foundations for developing critical thinking and ethical reflection; students in particular need frameworks for making sense of acquired knowledge through the exploration of ideas and ethical deliberation. An integrated humanities approach can blend the rigor of mathematics, science, and engineering with interpretive and contextual insights from humanistic inquiry. As part of this, STEM can be taught from a humanities perspective. For example, when students construct machine learning models, they can simultaneously examine how biased data perpetuates community harm; when studying computer science, they can engage substantively with questions of fairness and technology’s societal implications. By situating STEM learning within social, historical, and ethical contexts, education prepares students not merely as AI users but as critical thinkers capable of interrogating AI’s underlying assumptions while developing information and media literacy as essential competencies across all disciplines (Delphi panelists).

This integrated curriculum spanning all grade levels cultivates versatile individuals through systematic incorporation of literature, philosophy, linguistics, culture, and ethics into technical subjects. For example, a biology unit on biomedical engineering could integrate structured debates about genome editing with CRISPR or genetic modification of foods, wherein students apply classical philosophical frameworks—Stoicism, Aristotelianism, utilitarianism, nihilism, pragmatism, or existentialism—to evaluate the ethical dimensions of biotechnological possibilities (Delphi panelist).

**BOX 18**

## Strategies to support students' knowledge and understanding, through interdisciplinary frameworks

**Use philosophy to understand the difference between *can* and *should*.** Students learn what *can* be done through scientific method and subjects like math, science, and history. Equally important—especially in an AI-driven world—is grappling with what *should* be done. As Klein et al. observe, AI “models make words but people make meaning” (1). Philosophy, alongside other humanities disciplines, is uniquely positioned to guide this work. Integrating philosophical approaches fosters ethical, reflective, and deliberative thinking, with ripple effects that strengthen learning across subjects (Education Endowment Foundation 2014). Helping young people make sense of their knowledge in an AI world is essential.

**Multimodal approaches to learning.** AI has the potential to transform teaching and learning through multimodal approaches. While most instruction worldwide remains heavily text-based, research in the learning sciences, including from professor Heng Luo of Central China Normal University, shows that integrating text, images, speech, gesture, simulation, and embodied interaction significantly enhances comprehension, retention, and transfer (Luo 2023; Arifin et al. 2024). AI can facilitate these diverse ways of engaging with material, thus aligning with Universal Design for Learning principles that reduce learner variability by supporting diverse student needs (CAST, n.d.). Innovators from India to the U.S. are experimenting with AI-enhanced VR to bring biology and chemistry concepts to life (Hale et al. 2025; Luo et al. 2021; Rahimi et al. 2025; FotonVR 2025), while others provide historical figure chatbots for student interaction (Muncey 2025). Students, too, are finding creative applications, using AI to generate podcasts from class notes or create personalized quizzes for test preparation.

| To shift educational experiences in schools, **support students' social and civic development.**

An important purpose of schools in an AI-infused world is to help young people learn to live together. Across virtually every country in the world, schools are one of the most prevalent social institutions where children can meet others, in person, outside their immediate family and neighborhood. The role of schools in helping children develop empathy, respect, and tolerance, among other essential social capabilities, is an important counterbalance to the growing polarization of online discourse and interactions. Leaders worldwide have for many years argued that education should help people understand others, build a sense of shared humanity, and develop the skills to collectively achieve shared goals (Delors 1998). The actions in Box 19 can help advance this.

**BOX 19**

## Strategies to support students' social and civic development

**Youth-led civic learning in and outside of school.** Civic learning helps students develop the knowledge, beliefs, and behaviors needed to participate in civic life (Winthrop 2020). Schools and their communities can actively curate opportunities for civic learning that help young people develop civic knowledge and mindsets and practice civic behavior. This can include inviting students to make a contribution in a range of spaces in and out of school, such as identifying and solving a problem at school, making a piece of art for their neighborhood, speaking at a community event, participating in a discussion on difficult topics with neighbors, and joining a community improvement organization (Raine et al. 2025). Students can identify when AI extends their abilities—such as mapping local vegetation for a park project or practicing question-and-answer responses for a presentation. But AI cannot do these activities for students. Civic learning can also provide opportunities for students to work collaboratively, connect with others, and advance their social interaction and development.

**Civic games and simulations.** Well-designed games and simulations can be an effective strategy to support students' civic learning and foster social connection and development. These range from offline experiences like Model United Nations, a program schools around the world use to role-play United Nations-style debate and deliberation, to AI-enriched online experiences (Jesuit and Endless 2018). Game designers can collaborate with young people in harnessing AI to advance online civic learning games and simulations, including those currently used by educators such as online government role plays or games helping young people learn about relationships, communication, and collaborative action (Rivers and Bertoli 2024; iThrive Games 2025; European Union 2025).

To shift educational experiences in schools, **foster students' motivation and agency, including through real-world learning.**

Perhaps one of the most important roles of education systems in a world of rapid AI advancements is to foster students' internal motivation to learn alongside their ability to proactively navigate their own learning journey. One way to mitigate the potential risk of AI increasing students' passivity and self-efficacy is to encourage them to set and pursue a meaningful goal—that is, to foster their agency. Student motivation, engagement, and agency are interrelated, shaped by the environments they are in, and can change quickly with a shift in context (Reeve et al. 2022). Learning experiences that connect to students' interests and provide students with some measure of choice and control can help boost their motivation and engagement and give them a chance to develop agency (Anderson and Winthrop 2025). While there are many ways to boost student motivation, engagement, and agency—including the pedagogical approaches discussed above—experiential learning opportunities that help connect student learning to real-world issues and the world of work are especially well suited for doing this.

## BOX 20

# Strategies to support students' motivation and agency, including through real-world learning

**AI-supported experiential learning.** A wide range of teaching and learning approaches can help foster student motivation, engagement, and agency. These can include project-based and career-connected learning approaches. In these types of experiential learning experiences, AI can be a helpful tool for students: assisting in identifying problems in their school or community and developing and implementing a solution. Rather than offloading learning, students' use of AI in this scenario extends it—whether primary school students leverage AI to improve trash collection or secondary school students use it to develop a new product or business (Pandit et al. 2025; Riverside Learning Center 2025; Johansen 2018). Students can work together in groups on projects that are connected to local business, civil society, or governments. Online experiential learning simulations that are powered by AI have the potential to provide students opportunities to practice real-world skills from job interviews to pitch sessions for aspiring entrepreneurs (Mollick et al. 2024). This can provide students with opportunities to exercise their initiative and creativity in the context of real-world industries. Examples include entrepreneurship training programs around the world, like those supported by Junior Achievement, using AI agents to give secondary school students the opportunity to practice describing the product or service they have created as if they were pitching to an investor. Students receive real-time feedback and can repeatedly run through the simulation prior to presenting their ideas to business leaders in real life (JA Worldwide, n.d.).

**Student evaluation of AI.** Student agency can be further fostered by giving students the opportunity to evaluate their own use of AI. Students can run “light-touch” classroom audits, and document when AI improves productive struggle versus when it increases cognitive offloading. This information can then be shared with school leaders to inform policies and practices around AI use.

## 2 CO-CREATE EDUCATION AI TOOLS WITH EDUCATORS, STUDENTS, PARENTS, AND COMMUNITIES

AI companies whose products are designed wholly or in part for education must move beyond what one educational AI engineer describes as “superficial” consultations with educators, students, and families to deeply engage users in co-creating these products. This is especially true for companies that develop educator- and learner-facing tools. This collaborative approach supports educator and student buy-in and trust. It helps educational AI tools align with research-based learning practices that enhance rather than undermine essential teaching and learning experiences across diverse contexts and populations. This approach also addresses the persistent critique of many educational technology products: that they are simply old wine in new bottles—outdated pedagogical approaches in technically sleek packaging (Chen et al. 2025). By working together to co-create use cases and designs that leverage AI's potential for powerful

agentic teaching, these educational AI products can better serve students and their teachers. This can include the approaches below:

I To co-create AI tools, **establish teacher-tech co-design hubs**.

In interviews, teachers expressed frustration that “AI is being done to us, not with us,” and evidence suggests that successful AI companies meaningfully engage educators and other stakeholders from the earliest stages (Cukurova 2025). One approach that teacher organizations and unions, technology companies, governments, and researchers can take is to establish co-creation hubs in each major education jurisdiction so teachers (together with technologists) help lead the design of AI tools with the support of governments and researchers. The examples in Box 21 demonstrate the diversity of approaches these hubs can take.

## BOX 21

# Teacher-tech co-design hubs

**Netherlands National Education Lab AI (NOLAI).** NOLAI is a government-funded program that supports a university, teacher union, schools, and an ed-tech company collaboration on AI and education. Housed at Radboud Universiteit in the Netherlands, it brings together schools, teachers, ed-tech companies, and researchers to develop AI tools and use cases, test them in schools, and evaluate their impact (Radboud Universiteit 2025, Turner 2022). Key elements include:

- Government funding for a 10-year period (2022–2032) allows the lab to plan long-term and finance prototype educational AI products
- Joint purpose: All stakeholders collaborate to improve primary and secondary education identifying where AI can help
- Co-creation of projects: Research questions developed collaboratively between teachers, school leaders, researchers, and business
- Joint oversight through a steering group representing education, academia, and business, with government advice through a program council
- Evaluation and shared knowledge: academics are involved throughout, and findings will be shared with schools and businesses

**Public-private partnership for US National Academy for AI Instruction.** OpenAI, Microsoft, and Anthropic pledged \$23 million to the American Federation of Teachers to open a free training center, the National Academy for AI Instruction in New York City, for teachers and school staff to learn AI responsibly and ethically (Kelly 2025; Microsoft Source 2025a; Scragg 2025; Singer 2025c). Key elements of the center include:

- Joint leadership from union leaders and public/private stakeholders
- Free in-person and online training with pathways to credentials and continuing education credits
- Two-way communication enabling teachers to inform students how tech companies develop AI for students
- Tool agnostic: The Academy introduces multiple AI tools, including those from funders and other companies

| To co-create AI tools, **involve students and parents in AI decisionmaking and design.**

Education systems can involve students and parents in how to best use and roll out AI in education. In many classrooms, teachers are learning about AI use alongside students. Young people especially, if given the chance, have meaningful contributions to make on how AI is used well to support their learning, creativity, and connection (Rithm Project 2025; Center for Digital Thriving 2025). Some strategies for systematically partnering with students are included in Box 22.

## BOX 22

# Involving students and parents in AI decisionmaking and design

**Student AI Councils.** Creating AI councils within education jurisdictions can serve as mechanisms for embedding student voice in the co-design of AI tools to ensure their relevance, inclusivity, and pedagogical soundness before adoption. Council members could help develop AI policies for education jurisdictions or schools. Students on the council could also play an important role in vetting AI products before purchasing. Their participation can enhance alignment with real classroom needs and foster a sense of shared accountability. These councils could also provide early detection of misuse risks while cultivating AI literacy and civic engagement among youth. They could convene short design sprints where students collaborate with teachers, developers, and policymakers to test AI tools for usability, accessibility, and cultural context; incorporate student feedback into procurement and pilot decisions; and document lessons learned and share them through an open repository to guide future procurement and implementation of AI tools.

**Student and parent representation on technology committees.** Where schools and school districts have technology committees, student representatives from both the primary and secondary levels, as well as family members, should serve as members of these committees.

| To co-create AI tools, **support local language and community AI initiatives.**

Technology companies, governments, civil society organizations, researchers and communities can help include marginalized communities' and countries' voices in AI development, adoption, and integration. This is one important way to ensure that AI tools remain accessible, support digital equity, and do not perpetuate harmful biases. AI adoption must begin with listening; this means employing co-design processes with a focus on underserved communities, thereby allowing schools to develop applications reflecting their specific contexts, values, and educational needs rather than importing external solutions (Bozkurt et al. 2024; Weaver 2022; Ruiz et al. 2024; Delphi panelists).

Governments and technology companies can support civil society, researchers, and communities to develop a global network of local language initiatives to share information, resources, and advance the field. Important strategies to do this are investing in community-driven datasets, training multilingual models on representative samples, providing linguistic expertise for expanding machine translation

capabilities, and offering access to scholars who guide data collection, verify accuracy, and identify resources (Okolo and Tano 2024). Several initiatives across Africa demonstrate how this can be done, as seen in Box 23 below.

Additionally, technology companies themselves can play a critical role in ensuring greater linguistic equity. For example, BERT (Bidirectional Encoder Representations from Transformers) is a widely used open-source multilingual AI language model originally developed by Google. It understands words by examining their complete context—reading both what comes before and after each word simultaneously. This bidirectional approach makes BERT particularly effective at understanding meaning across different languages with varied grammatical structures and word orders. Several initiatives, some of which are profiled in Box 23, have recognized the potential of transformer-based architectures, such as BERT, to support African-language and Asian-language natural language processing, particularly in contexts where local language training data are scarce.

### BOX 23

## Local language and community AI initiatives

**Ghana Natural Language Processing (Ghana NLP) identified that most African languages are either absent from or poorly represented in systems such as Google Translate.** To address this gap, Ghana NLP developed ABENA (“A BERT Now in Akan”), a transformer model fine-tuned from multilingual BERT (mBERT) and adapted for Twi/Akan, a major Ghanaian language also spoken in Togo and Côte d’Ivoire. ABENA represents one of the first BERT-family models for a Ghanaian language, while Ghana NLP also focuses on improving datasets and methods optimized for low-resource settings (GhanaNLP 2025).

**Masakhane, a pan-Africa, community-driven research collective, promotes collaboration and capacity building in African language natural language processing.** Its community of researchers across the continent co-develops datasets, fine-tunes multilingual transformer models such as mBERT and XLM-RoBERTa, and publishes research that advances African-language processing (Masakhane, n.d.).

Beyond sub-Saharan Africa, regional transformer projects demonstrate the adaptability of BERT-family architectures to underrepresented languages. Two examples are IndicBERT for Indian languages and SEA-LION for Southeast Asian languages. These models, typically trained from scratch or fine-tuned from multilingual BERT, aim to produce tools that better capture linguistic and cultural nuances than English-centric models, thereby mitigating bias and improving inclusion in AI systems (Pava et al. 2025, 14).

Not all language initiatives need to use BERT-family architecture. Lesan.ai, based in Ethiopia, does not use the BERT encoder-only architecture, rather it applies machine translation techniques to expand access for underrepresented languages. Its system supports bidirectional translation among English and widespread Ethiopian and Eritrean languages such as Amharic and Tigrinya (Lesan.ai, n.d.).

### 3 USE AI TOOLS THAT TEACH, NOT TELL

LLMs typically operate through a simple prompt-response structure whereby the AI generates a complete answer in a single automated step (Wan et al. 2024), inducing passivity and cognitive offloading that undermines learning for many students. This design limitation is not inherent to all AI applications in education.

To use AI tools that teach, not tell, **technology companies can incorporate child-optimized approaches to their products.**

While many ed-tech companies may design AI tools drawing on vetted content and the science of learning and child development, general purpose AI technologies like LLMs are often not optimized to advance children’s learning. Fortunately, several technically feasible design strategies could help AI better support children’s development and learning. Technology companies, especially those with student-facing tools, can use the approaches in Box 24.

#### BOX 24

## Child-friendly product design

**“Antagonistic” design.** AI companies can design “antagonistic” AI tools and systems that compel users to confront their assumptions, build resilience, and develop healthier relational boundaries with AI technologies. These models may prove particularly beneficial in educational settings. Rather than consistently validating student choice, as current sycophantic AI tools often do, antagonistic models can be trained to challenge, critique, and productively disagree with users. This approach pushes students toward greater self-reflection and higher quality standards while strengthening their ideas and arguments and mitigating potential harm to their judgment, socialization, and resilience (Cai et al. 2024).

**Progressive disclosure.** AI models can be designed to utilize “progressive disclosure.” This technique, borrowed from User Design Experiences in online learning, involves showing only essential information and gradually introducing more complex features or explanations as the user progresses or requests them (UXPin 2025). Information is presented in smaller chunks so that complex information is more easily digestible and involves users in an iterative process of learning (Burns 2023b). If done well, this can help scaffold learning.

**Explainable AI (XAI).** Progressive disclosure can be combined with XAI techniques to help users understand how AI makes decisions (DiPaola et al. 2024). For example, a math tutoring platform using XAI would not simply give the answer “42” to the question, “What is  $15 + 27$ ?” Instead, the platform would explain its reasoning: “I broke this into parts:  $15 + 20 = 35$ , then  $35 + 7 = 42$ .” The student can then follow the steps and understand the method rather than just accepting the answer. XAI helps build what Turner et al. (2022) call “calibrated trust”—appropriate skepticism and reliance on AI-generated responses based on understanding rather than blind faith or complete rejection. When students can see how an AI system arrived at an answer, they can better assess its validity, identify potential biases, and make informed decisions about when to follow or override its suggestions based on the AI’s strengths and limitations.

**Cognitive Forcing Functions (CFF).** CFF strategies encourage analytical engagement with AI-generated content by requiring users to pause and reflect before receiving assistance (Buçinca et al. 2021). For example, a writing tutoring program might first ask students to identify three areas for improvement in their essay before providing feedback. This approach maintains students as the primary cognitive agents, enhancing their metacognitive and analytical skills through reflection rather than passive acceptance of AI-generated answers.

## 4 CONDUCT RESEARCH ON CHILDREN'S LEARNING AND DEVELOPMENT IN AN AI WORLD

There is an urgent need for high-quality research to track children's learning, well-being, and development in an AI world. As teaching and learning practices shift, including with AI use, the education community will need to regularly produce and use evidence on AI's influence on students' cognitive, affective, motivational, and emotional states. Research should be grounded in student experience, rigorously examine the potential risks and benefits of AI, and be shared in usable forms for practitioners and policymakers. When benefits and harms are identified, research must also illuminate the mechanisms driving those effects, moving beyond documenting outcomes to understanding why and how AI produces particular results for children (Deng et al 2025).

Research should focus on how to help students prosper in an age of AI, including real-time evidence development to help inform current AI practices before norms are entrenched. Below are some ways researchers and funders can help develop the knowledge base that can enable governments, education systems, and technology companies to mitigate AI's risks and maximize its benefits, including important areas to research.

To conduct research on children's learning and development, **prioritize research that surfaces AI risks and how to mitigate them.**

There is a need to carefully track and investigate the risks associated with AI, including those raised in this report across children's cognitive, emotional, and social development, including the risks of diminishing trust and increasing dependence. The education community needs to better understand the extent of these risks: how and under what conditions children are most susceptible. For example, further research is needed to understand how social chatbots can help individuals address social challenges and to identify the conditions that may lead to unhealthy engagement, such as overuse, dependence, or reduced human socialization (Franze et al. 2023). But most importantly, there is a need to better understand how to effectively mitigate these risks. What designs, levers, interventions, practices, norms, and policies are effective in reducing or eliminating the risks children face?

To conduct research on children's learning and development, **focus on research that understands how to leverage AI to benefit children's learning and development.**

Alongside better understanding the risks of AI and how to mitigate them, there is a need for research that uncovers the specific use cases, approaches, and contexts in which AI can be used to expand children's learning and development, including the benefits identified in this report. Which students

(e.g., neurodivergent or L2 learners) benefit from AI use? Under what conditions, with what designs, and through what types of implementation does AI help students expand their understanding, develop their skills, build their connections to others? For example, therapeutic interventions have advanced to a degree of sophistication whereby their responses are virtually indistinguishable from those written by humans (Zao-Sanders 2025; Heinz et al. 2025; Hatch et al. 2025). Research on how these increasingly available interventions could help address the youth mental health crisis students who lack regular or immediate access to mental health professionals is warranted (Heinz et al. 2025; Mahari and Pataranutaporn 2025).

| To conduct research on children’s learning and development, **include research that centers teachers.**

Teachers, whose experiences with AI have thus far been underexamined, should play an integral role in this research so the education community can gather a fuller sense of how AI impacts the entirety of the student learning process. Teachers are central to research on how to use AI—and what supports they need—to enrich student experiences and strengthen motivation, engagement, and agency. Researchers from the University of Oulu argue that teachers can also serve as models for AI training, feed data to the AI system (for example, lesson plan ideas), and determine assessment criteria (Celik et al. 2022). This participation also has the added effect of helping teachers better understand AI itself and its role in learning.

| To conduct research on children’s learning and development, **prioritize research on AI in low-resource settings, including examining system dynamics.**

To develop and apply effective AI in education in the most challenging environments, educators require context-specific research that examines AI interventions within specific contexts, cultures, and education systems. The effects of AI innovations are often incremental and may not completely visible by the end of a project or initiative. To truly capture the impact of AI on children’s learning and well-being, funders of AI initiatives can support research that extends beyond the lifespan of particular projects, especially when funding operates within fixed or predetermined time periods. This research must concentrate not only on the intervention itself but also on the systems and stakeholders that influence whether learning gains transfer and persist over time (Burns 2020, 50).

| To conduct research on children’s learning and development, **use a variety of research approaches.**

A variety of research approaches is needed, from real-time actionable insights to longitudinal studies, to address the range of questions on how to help children prosper in an age of AI. Box 25 outlines some of these approaches.

**BOX 25**

## Diverse research approaches to study AI and education

**Produce practical, usable research grounded in students, teachers, and parents' lived realities.**

Educators and AI practitioners need research to inform their decisions, yet may often eschew it because they view research as lacking practical utility and disconnected from the lived experiences of those implementing and receiving interventions. To address these issues, researchers can embrace a “utilization-focused” approach wherein research is designed for specific users rather than a general audience, and primary intended users are involved in the research design and process (Patton 2008). By including student, teacher, and parent perspectives, expertise, and experiences, research findings produce deeper, more nuanced insights than external studies alone, leading to more compelling and coherent explanatory narratives. This ultimately makes the research more evidence-based, humanistic, and usable to decisionmakers, practitioners, and educators (Burns 2020 50–51).

**Rigorously assess AI's impact on children now through immediate prospective research.** The social media experience demonstrates why urgency is critical. As Bernstein (2022) notes, by the time research on social media's effects on children became widely available, platform business models had already been built around engagement-maximizing algorithms, users had developed dependency patterns, billions in revenue streams were at stake, and regulatory efforts faced entrenched resistance and significantly higher implementation costs. The AI landscape still remains malleable. Companies with products children access continue experimenting with design features, users have not yet formed fixed usage patterns, and fundamental questions about AI's impacts on children remain open (Bernstein 2022). Researchers must document and disseminate what is known about AI's effects on children's learning and well-being now, while evidence can still inform design decisions, policy frameworks, and educational practices before harmful patterns become embedded in both technology and behavior. The window for preventive action is narrow (Bernstein 2022, 2025). Children are in critical developmental periods where cognitive, social, and emotional capacities are being formed. Interventions that shape AI's role in education today can prevent the need for far more difficult and costly remediation efforts later.

**Examine whether general-purpose and companion AI platforms should be researched and regulated similarly to drug trials.** Many prescription drugs are available and widely used even though researchers and doctors may not fully understand the exact mechanisms by which they produce their effects. Governments might consider whether AI technologies used by virtually all primary and secondary students should be studied in similar ways to pharmaceuticals, that is, through a series of clinical trials that produce risk-benefit analyses and public labeling of proven harms to student learning or well-being. As part of this, researchers and education systems could continuously monitor these applications to assess whether new warnings should be added or whether access should be limited for students of certain age groups or with particular neurological or affective conditions (Delphi panelist).

**Create evidence-based pilot programs and impact demonstrations.** Greater research on AI in education can generate opportunities to develop pilots, sandboxes, proofs-of-concept, and action research initiatives in partnership with research institutions. Such efforts can demonstrate measurable impact, accelerate the translation of evidence into practice, and determine which models of AI work under

what conditions (Joel Mitchell, personal communication, August 29, 2025). These pilots should aim to identify the tipping point at which investment in AI leads to measurable positive impacts on problems facing low-resource countries. Educators and policymakers in sub-Saharan Africa interviewed for this study emphasized the importance of demonstrating how strategic AI investments can address systemic educational challenges such as teacher shortages, infrastructure gaps, and inadequate learning tools across different cohorts and demographics. Beginning iteratively with small-scale approaches may provide essential proofs of concept before scaling. Establishing mechanisms for systematic evaluation and evidence-based expansion enables the creation of frameworks that balance innovation with prudent resource management, creating pathways from promising pilots to sustainable, large-scale AI initiatives in education.

**Produce rigorous, empirical, and longitudinal research on AI in education.** AI is in urgent need of rigorous and longitudinal research. Indeed, most educators are implementing practices without the benefit of empirical evidence. Such research can involve properly isolating variables, measuring actual learning outcomes, considering long-term impacts, exploring innovative instructional designs, and providing sufficient detail for replication (Trabelsi 2025).

**Consider the Bradford Hill criteria for student emotional well-being research.** Researchers can apply the Bradford Hill criteria to study the impact of AI on students' emotional well-being. These widely respected scientific principles assess the strength of evidence for causal relationships between variables. They are particularly valuable for determining whether exposure to a risk factor, such as an AI companion bot, determines adverse outcomes, especially in contexts in which randomized controlled trials are difficult or impossible to implement (Lembke 2023).

## Prepare

Governments, education systems, funders, civil society organizations, families, communities, and young people themselves must collaborate to help students navigate an AI world. Preparation requires building the knowledge, capacity, and structures for ethical and effective AI integration, ensuring that schools develop clear AI visions with dedicated resources, organized adoption processes, and measurable evaluation criteria to track implementation success.

To ensure that young people and our education systems are prepared for an AI world, governments, funders, private sector, education systems, families, and communities can work together to advance these four main actions:

5. **Promote holistic AI literacy for students, teachers, parents, and education leaders.** Integrating ethical understanding, critical thinking, and practical skills across curricula and communities helps to ensure all learners can engage with AI confidently, creatively, and responsibly.
6. **Prepare teachers to teach with and through AI.** Both pre-service preparation and in-service training equip future and current teachers, respectively, with the knowledge, skills, and confidence to use AI responsibly, model critical thinking, and preserve authentic student learning.

7. **Provide a clear vision for ethical AI use that centers human agency.** Education systems should prioritize student agency, ethical learning, critical thinking, and ensure that AI tools are vetted against rigorous criteria that promote human reflection, transparency, and genuine intellectual growth.
8. **Employ innovative financing strategies to close the AI divide.** There are multiple ways that actors, especially governments and education systems, can help reverse the growing AI divide, from equity-focused policies to innovative financing mechanisms for expanding school connectivity.

## **5 PROMOTE HOLISTIC AI LITERACY FOR STUDENTS, TEACHERS, PARENTS, AND EDUCATION LEADERS**

One of the most important steps to prepare education systems to harness AI in a way that will help learners prosper is to cultivate holistic AI literacy skills of the adults and children within education systems. Developing AI literacy skills does not require teachers or students to spend extensive time on screens. Rather, it can include in-person discussion and reflection alongside some practice and exploration of AI tools themselves.

TeachAI, a global coalition of educationalists, defines AI literacy as the knowledge, skills, and attitudes associated with how AI works, including its principles, concepts, and limitations, as well as how to use AI (TeachAI 2025). Participants in this study emphasized the need for “holistic” AI literacy, which empowers learners to “engage with, manage, create, and design AI systems while understanding the ethical, social, and cognitive implications of their use” (European Commission and OECD 2025). With this framework in mind, governments and education systems can provide guidance to educators on what students should know and be able to do in order to be considered AI literate.

| To promote holistic AI literacy, **adopt holistic AI frameworks.**

There is no shortage of holistic AI frameworks that governments and education systems can draw upon. Education systems and governments can look to other countries who are embracing holistic AI approaches and rolling them out nationally. For example, the government of China requires that students receive at least eight hours of AI education a year beginning at age six, outlines four key areas of AI literacy (cognition, skills, thinking, and values), with students in lower grades focused on awareness of AI while students in higher grades focus on creating and problem-solving with AI (Australian Government Department of Education 2025).

There are also a wide range of useful holistic AI literacy frameworks developed by global networks, universities, and civil society organizations that are used by education jurisdictions around the world. The AI Literacy Framework for Primary and Secondary Education, developed jointly by the Organisation for Economic Co-operation (OECD) and European Commission with support from Code.org and an expert group of international educators, researchers, and technologists, approaches AI literacy broadly; as cited above this framework includes learning how AI works, from understanding the ethics to how to harness its power in problem solving (European Commission and OECD 2025). There are multiple teacher professional development offerings; The International Society for Technology in Education and Association for Supervision and Curriculum Development (ISTE+ASCD) offers AI courses, and UNESCO’s Guidance on Generative AI outlines holistic teacher competencies (ISTE 2025; Miao and Holmes 2023).

Annex A reviews a range of AI literacy frameworks from China and Uruguay to Digital Promise and MIT.

| To promote holistic AI literacy, **create or adopt guidelines for AI literacy.**

Education systems can prioritize holistic AI literacy guidance. While many countries operate highly decentralized educational systems, both centralized and decentralized approaches can successfully implement AI literacy frameworks. Countries with centralized systems, such as China and Estonia, have implemented comprehensive national AI literacy guidelines that provide consistent standards across all schools (Asian College of Teachers 2025, e-Estonia 2025). In Singapore, teaching on AI's uses, risks, limitations, and ethical considerations is embedded across different school subjects, along with instruction on verifying online information and understanding data security and privacy (Ang 2024).

Countries with decentralized education systems have also established effective national frameworks. India, despite significant state-level autonomy in education, developed national AI literacy guidelines through the National Education Policy 2020 and the Central Board of Secondary Education, which introduced AI as a subject for students in grades 9–12 and created an AI Facilitator Handbook for educators (Ahuja et al., n.d.; Mehra 2020).

There remains a compelling advantage to establishing centralized national guidelines that define AI literacy. Without coordinated national involvement, students may receive vastly different levels of AI literacy preparation depending solely on their geographic location, creating inequitable educational outcomes in an increasingly AI-dependent society. National guidelines can ensure consistent foundational knowledge while still allowing for local adaptation to meet specific regional needs and contexts.

| To promote holistic AI literacy, **support systemic AI literacy approaches.**

In the interest of deep AI literacy, education systems can integrate these approaches across grades/levels and the curriculum, rather than confining it to isolated courses or particular disciplines such as computer science (DiPaola et al. 2024, 8; European Commission and OECD 2025; Dhar 2025). A growing number of secondary schools and higher education institutions are adopting this comprehensive approach, which helps students encounter these concepts throughout their learning experiences.

Implementation can reflect developmental appropriateness, beginning with playful and offline learning opportunities for young children that introduce foundational concepts through age-appropriate experiences. As students mature, they can progressively experiment with AI systems, building both technical familiarity and critical understanding through hands-on engagement. This scaffolded approach recognizes that meaningful AI literacy develops over time and requires repeated encounters across contexts.

| To promote holistic AI literacy, **support peer-to-peer AI literacy.**

Students can play a powerful role in AI literacy with their peers. Education systems can train student facilitators to develop and deliver short sessions to their peers on a range of topics related to AI use, their education, and well-being. For example, sessions could address prompt hygiene, verification habits, when not to use AI, debates on ethical uses of AI, and strategies for switching from AI to human help. Peer-led AI literacy normalizes responsible, inquiry-based use of generative tools while easing teacher capacity gaps. It transforms AI from a passive technology into shared learning grounded in reflection and accountability.

| To promote holistic AI literacy, **include families and communities in AI literacy.**

Education systems can work with families and caregivers to extend AI literacy skills and guidance from school into home environments. This requires training school administrators and teachers in effective family engagement approaches that center families as active partners and provide space to co-create solutions (Winthrop et al. 2021). Through intentional communication, trust building, and understanding of diverse belief systems, teachers can learn about caregivers' perceptions of children's technology use and establish shared guidelines for responsible AI engagement. Such partnerships can create consistent messaging about healthy relationships with technology and authentic human connection (Delphi panelist).

AI literacy will be especially useful if it extends beyond students to encompass the entire school community. Principals, teachers, students, and parents can all develop an understanding of how algorithms shape their daily lives and acquire the capacity to critically interrogate the systems operating around them (Dhar 2025). This shared literacy creates a foundation for informed decisionmaking about AI integration while fostering the collective critical consciousness necessary to navigate an increasingly algorithm-mediated educational landscape. India's "AI Samarth" initiative exemplifies this approach by providing AI literacy for students, teachers, and parents, thereby fostering a greater sense of shared understanding and alignment around AI and education across the entire school community (Central Square Foundation 2025).

## **6** PREPARE TEACHERS TO TEACH WITH AND THROUGH AI

As ISTE correctly notes, "there are no AI literate students without AI literate teachers" (ISTE 2024). The integration of AI into education demands a fundamental transformation in how we prepare and support educators. Teacher pre-service or preparation programs globally have historically lagged in modeling effective technology integration (Burns 2023a). Teachers require comprehensive guidance in order to navigate the proliferation of AI education tools, develop expertise in productive and ethical AI integration, and master the delicate balance between leveraging AI's capabilities and preparing students to use these tools while preserving authentic student learning and agency. Education systems and civil society can support this transformation across pre-service preparation—the education and preparation a person receives *before* becoming a teacher—and in-service professional development—the support and training *current* teachers receive.

| To prepare teachers, **integrate AI literacy into pre-service teacher preparation.**

Teacher pre-service preparation programs can integrate AI literacy throughout their curricula through the following approaches:

- Provide practical experience integrating AI into classroom instruction rather than limiting preparation to theoretical AI literacy knowledge, enabling future teachers to develop essential confidence and self-efficacy in AI implementation (Park 2023)
- Develop comprehensive understanding of AI's dual nature in education by helping pre-service teachers recognize how AI can both support and potentially harm children's learning and development, while building a repertoire of use cases and instructional activities that strengthen students' abilities to perceive, remember, form concepts, solve problems, imagine, and reason (American Psychological Association 2025a)

- Emphasize metacognitive skill development to prepare future teachers to cultivate students' awareness of when and how to use AI tools responsibly without compromising students' cognitive development (Bozkurt et al. 2024)
- Establish new graduation requirements and instructional models that require faculty in teacher training institutions to develop their own capacity for modeling ethical, responsible, and effective AI integration, addressing current uncertainties about appropriate AI use in teaching and learning (Weiner et al. 2024; Burns 2023a)

To prepare teachers, **support teachers in teaching with AI through in-service professional development.**

In-service training for practicing teachers, also known as continuing professional development, can also support teachers in teaching with AI and can be strategically designed and implemented through:

- Differentiated and sustained professional development that precedes and accompanies AI implementation. This professional development should strengthen teachers' content knowledge so they can verify AI outputs. It should also enhance teachers' pedagogical content knowledge to seamlessly blend content, teaching strategies, and AI technology into high-quality learning experiences.
- Instructional approaches that help teachers balance and navigate the complex relationship between AI use and pedagogical choices. Teachers need to recognize that different instructional methods create varied opportunities for cognitive offloading. They may need guidance on selecting appropriate strategies and titrating AI use to maximize learning benefits across diverse pedagogical approaches that include both direct instruction and learner-centered methodologies (Hill et al. 2022).
- Curriculum-aligned activity design that employs AI meaningfully while preventing opportunities for students to offload cognitive work or replace genuine effort with AI-generated outputs.
- Collective professional learning communities organized by subject area and grade level, ensuring professional development remains consistent with standards and policies, grounded in current research on effective AI use, and supported by ongoing follow-up to improve implementation fidelity and transfer of learning (Burns 2023a).

A number of education systems are incorporating AI into teacher training curricula. In Singapore, for example, by 2026, the National Institute of Education will offer training in AI in education to all undergraduate, post-graduate, and in-service teachers (Kai 2022).

Students can also be involved in teacher professional development. For example, training programs can develop student mini-labs, where teachers bring draft activities and students test them for clarity, temptation to over-offload, and inclusion. This provides teachers with rapid feedback before lessons go live.

## **7 PROVIDE A CLEAR VISION FOR ETHICAL AI USE THAT CENTERS HUMAN AGENCY**

Thinking and learning are not tasks to be outsourced to AI—they are how students build their identity, agency, and dreams. While AI offers speed and fluency, it cannot experience wonder, wrestle with doubt, or choose values. Education systems can help students see that struggling to find their own words and

pursue their own questions is what makes learning meaningful and life worth living (Delphi panelist). They can do this by developing a clear vision for how AI can be used to help ethically advance human agency.

To provide a clear vision for ethical AI use that centers human agency, **first create a solid policy.**

Education systems from the school level on up can develop and broadly disseminate this vision either through creating new policies or incorporating the vision into existing relevant guidance and policies. Educational policies can encourage students to use AI to expand their ability to explore and learn—processes that require critical thinking and personal engagement rather than cognitive offloading. For example, in the U.S., the Washington Office of Superintendent of Public Instruction has issued guidance on using AI in K–12 schools that promotes a human-centered approach. This guidance states that good AI use “always starts with human inputs and inquiry, and always concludes with human reflection and edits.” It notes that AI should be used to empower students to actively participate in their education, not replace student development (Washington Office of the Superintendent of Public Instruction 2024, 14). Studies show that when schools or universities have AI policies, teachers and students are more likely to use AI (Gallup 2025; Xiao et al. 2023).

## **8 EMPLOY INNOVATIVE FINANCING STRATEGIES TO CLOSE THE AI DIVIDE**

Designing for equity is often a political choice involving prioritizing access to quality learning, deliberately directing resources to marginalized communities, and developing locally tailored initiatives. To do this, governments, educational systems, funders, and technology companies can embrace inclusion, justice, and equity, embodying what Leandro Folgar Ruétalo, vice president of innovation at Universidad Católica del Uruguay, describes as “equity extends to all human beings.” This ensures that all students can meaningfully participate in AI transformation. Rather than using readiness criteria that exclude marginalized communities, this approach requires supporting underresourced education systems, schools, and communities through targeted policies and strategic resource reallocation (Delphi panelist). There are a range of strategies that can help advance this priority.

To employ innovative financing strategies to close the AI divide, **promote equitable access to AI infrastructure through innovative financing for education.**

Many education systems lack the funds or funding to adopt and implement AI in education initiatives. Given the closure of and cuts to many bilateral aid agencies, governments in the Global South find themselves even further behind in terms of funding and traditional funding models may be insufficient for AI and digital infrastructure transformation.

These governments will need to explore different options to diversify AI funding, which may include innovative approaches that create sustainable educational infrastructure while building local capacity and ownership. One strategy for achieving this diversification is to draw on financing mechanisms already used in broader education contexts (Patrinos and Tanaka 2024). Innovative Financing for Education (IFE) is an umbrella term that encompasses new sources of funding, new actors, and new ways of sharing costs, risks, and responsibilities between private and public actors to provide more funding for education. IFE involves a number of alternative financing methods, many of them grounded in market-based and private sector mechanisms (Avelar et al. 2020, 4). These approaches can be adapted to support AI infrastructure in education.

## IFE for AI infrastructure

**Public-private partnerships** (PPPs) are formal arrangements between government and private organizations to jointly finance and deliver services. Technology companies have a long history of such partnerships with governments. One of the most well-known educational technology PPP may be Intel Teach, through which Intel trained millions of teachers globally to integrate technology into lessons and cultivate students' critical thinking, problem-solving, and collaboration skills.

These PPPs continue in the age of AI, with companies voluntarily providing money, tools, and services to schools (Singer 2025c). For example, Google committed \$1 billion to AI education, digital well-being, job training, and AI research for U.S. students. The company is making advanced AI tools, including a new Guided Learning mode, free to students and providing U.S. secondary schools with Gemini For Education (Pichai 2025). It plans to expand its AI for Education Accelerator to over 100 colleges and universities and create an online hub for AI education resources. In India, OpenAI partnered with the Association for Reinventing School Education to introduce AI tools and training to K–12 educators (Economic Times 2025). These PPPs must be carefully structured so they are not simply marketing and customer capture strategies but rather prioritize authentic educational objectives developing learners' critical thinking about AIs (Singer 2025a).

Such partnerships can extend beyond tech companies to include schools, pairing, for example, well-resourced international or private schools in capital or large urban centers with local low-resource government schools. In such arrangements, government school teachers work alongside or receive coaching from private school teachers on AI use and private schools may make computer labs or laptop carts available to students from government schools as part of after-school programs. Teachers from low-resource schools could adapt AI tools for offline use and share their adaptations through open toolkits.

**Blended financing mechanisms** combine government funds, private investment, philanthropic donations, and development aid to leverage each capital type and reduce individual risk. This mechanism attracts new sources of private capital directed towards initiatives that advance development goals, mitigate risk, and generate potential financial returns.

**Build-Operate-Transfer (BOT) models** involve private entities in constructing and operating educational infrastructure before transferring ownership to the public sector. As the name suggests, there are three parts to a BOT: **Build** (design and develop the infrastructure), **Operate** (run, maintain, and monitor the system), and **Transfer** (after a fixed period, typically 30 years, transfer ownership so that the host can sustain without outsourcing).

BOT models have been successfully implemented in other government sectors and public infrastructure projects (Patrinos and Tanaka 2024). However, it is difficult to find examples of BOTs involving educational technology—typically, one of the letters, usually “T” for transfer, seems to be missing. Nonetheless, sub-Saharan African policymakers interviewed for this study advocated for BOTs as a possible infrastructure provision model. This approach allows educational institutions to access immediate infrastructure improvements while spreading costs over time and ultimately gaining full

ownership without requiring substantial upfront capital investment, thus moving beyond traditional “charity-based” interventions toward sustainable business models that create mutual value for private partners and educational institutions.

**Results-Based Financing** (RBF) links payments to verified achievement of measurable outcomes rather than funding activities or inputs. In the case of social and development impact bonds, one form of RBF, social investors provide up-front capital and assume financial risk traditionally borne by governments. By making disbursements contingent on predefined targets, education stakeholders can be incentivized and supported to strengthen systems that support the achievement of learning outcomes (Gustafsson-Wright et al. 2017). Though seemingly not common in terms of educational technology, RBF approaches are increasingly common in education.

One example of an RBF is the Sierra Leone Education Innovation Challenge (SLEIC), a three-year \$18 million partnership between the Government of Sierra Leone and the Education Outcomes Fund, with implementing partners such as Rising Academies, Save the Children, and Street Child working across 325 public primary schools with approximately 134,000 children to boost literacy and numeracy outcomes.

SLEIC exemplifies results-based financing (also called outcomes-based financing), because payments are contingent upon attaining specified results rather than simply delivering predetermined inputs or activities and the financing structure ensures that money is spent only on activities proven to have an impact. The model explicitly encourages partners to innovate—several, for example, use digital tools such as WhatsApp and SMS/phones for teacher coaching. It also employs rigorous independent evaluations and caps per-child costs at \$36 to support future scalability. Results published for Year 2 show statistically significant learning gains in math and English, with payments tied to these verified learning outcomes (Education Outcomes Fund 2024).

To employ innovative financing strategies to close the AI divide, **use regulatory license requirements to drive private sector support for school connectivity.**

Governments can use creative license requirements to support bridging the AI divide. South Africa exemplifies this by strategically leveraging its regulatory authority to mandate enhanced educational connectivity through telecommunications licensing requirements and spectrum allocation conditions, making school Internet access a fundamental prerequisite for operators to maintain their licenses. It accomplishes this through two main mechanisms:

- **Mandatory School Discounts (E-rate):** All licensed Internet providers must offer schools at least a 50% discount from regular Internet prices. This isn't optional; it is a legal requirement that covers everything from Internet access to equipment and phone calls to tech support. Companies that fail to comply can be fined by the South African government (Independent Communications Authority of South Africa, n.d.; Vecchiatto 2009).
- **Connection Quotas (Universal Service Obligations):** In exchange for operating licenses, or in order to buy radio spectrum from the government, all telecommunications companies are assigned specific numbers of schools they must connect to the Internet. Recent spectrum license holders must connect over 18,500 public schools nationwide along with hospitals, clinics, and police stations within a

certain radius of schools. These renewable licenses include specific metrics for school infrastructure development, evolving from basic computer provision to comprehensive digital infrastructure and connectivity (Independent Communications Authority of South Africa 2021; Ellipsis 2025).

Governments can go even farther and mandate that Internet and mobile service providers offer simple, comprehensible connectivity plans for education systems or provide credits to teachers and students individually at fixed or subsidized rates. These requirements will enable people to access basic connectivity options that support their educational goals with ease and at an affordable cost (West 2025).

To employ innovative financing strategies to close the AI divide, **develop options for equitable distribution of AI tools.**

Governments can also help reverse the large and growing AI divide by pursuing creative pathways for deploying AI to marginalized communities.

## BOX 27

# Strategies for equitable distribution of AI tools

**Strategic Infrastructure Deployment:** Governments can consider deploying tiered infrastructure solutions that accommodate varying technological capacities within education systems. This might include offline or low-bandwidth AI tools, such as local language models embedded directly within school devices, so that technical limitations do not become barriers to educational access and opportunity (Delphi panelist).

**Government-developed AI platforms:** To limit both reliance on for-profit technology companies and save money in licensing fees, governments can build their own AI platforms and tools for education systems, as the government of Indonesia has done. In South Africa, the Department of Basic Education has overseen the development of digital content for South Africa's national curriculum. The content is completely accessible, is hosted on Department of Basic Education servers ("DBE Cloud"), is device-agnostic (mobile, tablet, and computer), and can be accessed for the equivalent of two South African cents per day (Burns and Santally 2019, 51).

**Open-source AI models:** School districts in resource-constrained contexts can embrace open source AI tools to scale equitable and quality access to advanced AI capabilities, provided appropriate technical support structures are in place. South Africa can serve as an exemplar in how to leverage open-source technologies, given its established leadership in open education through institutions such as the South African Institute for Distance Education and successful initiatives including Siyavula and OER Africa (Burns and Santally 2019). Open-source libraries such as TensorFlow, PyTorch, and Hugging Face Transformers, along with repositories like GitHub, provide accessible resources for implementing these AI capabilities.

**Deploy alternative mechanisms for expanding Internet access:** While regions like sub-Saharan Africa are achieving greater and more affordable Internet access, there remains a persistent and multifaceted

(geographic, economic, and gender-based) Internet divide. To address this challenge, governments can invest in both policy and technical solutions. For example, governments can lease unused fiber optic cable or “dark fiber,” invest in resource virtualization or satellite-based Internet, build government-owned public education networks, levy a tax to create Universal Service Funds, and/or use “white spaces”—wireless technology that utilizes unused television and radio frequencies to provide wireless broadband connections. Countries such as Malawi, South Africa, Namibia, Nigeria, Tanzania, and Zambia have expanded Internet access through various combinations of these initiatives (Burns and Santally 2019).

To employ innovative financing strategies to close the AI divide, **prioritize AI access for girls and other marginalized groups.**

The AI divide is multifaceted, with girls particularly disadvantaged in terms of AI use (Economist 2024). Technology companies, funders and governments can begin to address this divide by funding implementation of the nonbinding Global Digital Compact (GDC), a comprehensive multi stakeholder framework adopted in 2024 that establishes principles and actions for an open, secure digital future. The GDC commits governments and non-state actors to bridging the digital divide, fostering data privacy, combating misinformation, and establishing international AI governance.

The GDC calls for gender equality and the meaningful participation of all women and girls in digital spaces to attain the United Nations’ Sustainable Development Goal 5 (achieving gender equality and empowering all women and girls). It advocates eliminating barriers impeding girls’ access to technology and AI, greater female leadership in technology decisions, mainstreaming gender perspectives into digital connectivity strategies, countering technology-facilitated gender-based violence, and developing data standards designed to prevent bias, discrimination, and human rights violations.

## Protect

Governments, education systems, technology companies, families, and students themselves share responsibility for ensuring that the technology children use daily protects their educational development and well-being. Protecting means safeguarding students’ privacy, safety, emotional well-being, and cognitive and social development through ethical AI tools that enhance learning while fostering emotional development, strong relationships, and productive struggle.

This requires technology companies to create safe, age-appropriate experiences. It requires governments and education systems to promote AI’s benefits while mitigating harms through enacting reasonable guardrails. It also requires that the adults in students’ lives—families, educators, community leaders—work alongside students to establish structures, routines, and modeling of ethical and responsible AI use so that students develop the habits of mind and discernment necessary to thrive in an AI-infused world. Multiple stakeholders can work together to:

- 9. Break the engagement addiction and design platforms that are centered around positive mental health for children and youth.** AI platforms should be designed to minimize manipulative engagement

features and actively promote positive mental health and well-being for young users. This includes shifting measures of success for their products.

10. **Establish comprehensive regulatory frameworks for educational AI.** Effective regulatory frameworks should be flexible and should align AI governance with student rights and safety, embedding accountability, ethical design, and data protections across all stages of educational AI development and deployment.
11. **Procure technology that protects students' privacy, safety, and security.** Education systems can safeguard student well-being by leveraging their purchasing power to procure only technologies with built-in privacy, safety, and security protections that prevent misuse of student data and ensure responsible AI use in schools.
12. **Support families to manage children's AI use at home.** Families and students themselves can play an important role in mitigating AI overuse at home. This can include conversations between students and their peers and families on their AI use and implications for their learning and well-being as well as providing families with quality information on AI to guide their decisions at home.

## **9 BREAK THE ENGAGEMENT ADDICTION AND DESIGN PLATFORMS THAT ARE CENTERED AROUND POSITIVE MENTAL HEALTH FOR CHILDREN AND YOUTH**

*The question isn't whether technology will reshape society—it is what incentives will drive that transformation.*

— Center for Humane Technology, 2025

As AI capabilities rapidly advance, developers of student-facing AI systems can prioritize ethical considerations and implement comprehensive strategies that safeguard student well-being while harnessing educational benefits.

**To break the engagement addiction, require online products to meet safety standards.**

Governments can require online services that children access, including AI, to meet standards that protect children's safety and well-being. Technology companies designing AI products used by children can follow guidelines put out by child development specialists, such as the "Children and AI Design Code" developed by the 5 Rights Foundation. These and other groups describe a clear distinction between products developed for children and those for adults. The American Psychological Association warns that manipulative design can interfere with "the development of healthy real-world relationships" (American Psychological Association 2025a). Governments can require technology companies which students access to develop products for children that meet criteria like those put forward in the Children and AI Design Code below.

## AI safety standards for children

1. **Developmentally appropriate.** The AI system is designed and operated to account for children's differing needs and vulnerabilities at different ages and stages of development by design and default.
2. **Lawful.** The AI system complies with applicable local, national, regional, and international law, rules, and regulations across all domains, including, but not limited to, children's rights, data protection and privacy, child exploitation and abuse, illegal and harmful online content, anti-discrimination laws, consumer protection, intellectual property, health and safety, and education.
3. **Safe.** The AI system does not create or amplify risks to the wellbeing or the physical, mental, and emotional safety of children, including privacy and security risks.
4. **Fair.** The AI system treats children and their data fairly and creates outcomes that are just and equitable for children.
5. **Reliable.** The AI system functions as expected. Performance and outcomes remain robust over time, including in unexpected or harsh conditions, or when atypical data is introduced....Humans can intervene to take control if required.
6. **Provide redress.** It is easy for children and those who represent their interests to report concerns and to seek actionable and effective recourse and remedy...It is easy to appeal both in-app and without logging in. Emerging concerns or extreme incidents are swiftly and effectively addressed, including by providing an easily reachable human contact.
7. **Transparent.** Stakeholders (including children) have access to adequate and accessible information to have a reasonable understanding of what the AI system does, its impacts, the measures taken to account for the capacities and needs of children, and the efficiency of these measures...
8. **Accountable.** A continuous chain of human and organisational responsibility is established across the whole AI system's value chain and lifecycle...
9. **Uphold rights.** The AI system upholds children's rights...[as per the United Nations Convention on the Rights of the Child]...including their right to life, to participate, and to protection. Inherent in children's right to life is their right to be fully realised as individuals, including meeting their need for agency, connectedness, and purpose. The AI system prioritises children's best interests and takes account of their voices and opinions (5 Rights Foundation 2025, 21-22).

| To break the engagement addiction, **stress test AI platforms for safety.**

One important way technology companies can help ensure AI products and systems are designed to protect children's safety and privacy is to engage in robust testing prior to release. Such testing can include the approaches in the box below.

## BOX 29

# Strategies to stress test AI platforms for safety

**User testing:** AI systems can undergo thorough and continuous user testing to determine how actual users interact with a system under normal conditions. User testing examines usability, functionality, user experience, and whether the system meets intended use cases. It can identify and mitigate potential unintended consequences and negative impacts before widespread release.

**Technical or ethical checks:** This may include limiting inappropriate content exposure, disallowing harmful academic or personal suggestions, enforcing usage breaks to prevent over-dependence, monitoring for signs of social withdrawal, and ensuring AI interactions support rather than replace human relationships and learning experiences (Mahari and Pataranutaporn 2025).

**Inclusion of safety features:** Such features could include parental controls that allow parents to link to their children's AI accounts, set limits on time and what children can see, and receive notification if their child's online behavior suggests potential self-harm (OpenAI 2025c).

**Red teaming:** Red teaming is the comprehensive, methodical probing of an AI system through sustained, coordinated efforts by skilled professionals who simulate real-world adversarial attacks using the full spectrum of techniques available to actual attackers. Once they identify these vulnerabilities, they go about fixing them (Hao 2025).

To break the engagement addiction, **employ advisory boards to ensure safeguards are in place to protect children from harms.**

Technology companies whose products are used by students can create advisory boards that include "scientists, youth, ethicists, health professionals, and other stakeholders who are primarily charged with the protection of adolescents" (American Psychological Association 2025a).

To break the engagement addiction, **engage students in developing safeguards.**

Students can co-draft default settings for time limits, logging, and nudges that push learners back to human help when tasks cross a difficulty threshold. Students can participate in safeguarding privacy, safety, and data ethics when schools procure or deploy AI tools by becoming "privacy stewards." Such participation can build institutional transparency and trust, help schools translate technical policy into accessible language, empower students as informed digital citizens, and strengthen compliance through bottom-up accountability. To do this, schools can first select and train a small team of student "privacy stewards" in basic data rights, vendor transparency standards, and digital wellbeing principles. These stewards would then be provided with vendor one-page checklists summarizing district or school practices on data use, retention, and sharing. Student privacy stewards would beta-test AI applications

under consideration for purchase and flag privacy concerns or questions for district purchasing officials for consideration.

| To break the engagement addiction, **shift financial incentives from engagement to utility.**

Investors and funders can help break the engagement addiction by shifting how they measure success. Current financial incentives for child-facing technology platforms often reward companies for time spent on the platform regardless of what young people are doing (Sindermann et al. 2024). For example, in 2022, social media companies earned \$11 billion from advertisements to children under eighteen (Raffoul et al. 2023), tracking clicks, time spent, and advertisement impressions—metrics focused on “engagement” rather than impact on users. From a profit-maximizing perspective, it makes little difference whether a teenager spends an hour on a social media or AI platform learning cooking skills or an hour learning about self-harm (Mishra and Warr 2025).

This “time on platform” metric extends beyond social media to AI companions and educational technology. AI friends, whether embedded in social media or standalone products, are designed to maximize engagement rather than quality of experience. Even ed-tech products face these perverse incentives. One ed-tech company leader reported that integrating AI into their product nearly halved the time teachers needed to generate administrative reports—a clear win for educators—yet faced board criticism for declining user engagement. Seth Reynolds, an education leader at EY-Parthenon, argues that traditional ed-tech success measures must be updated to include metrics that assess impact (Seth Reynolds, personal communication, September 22, 2025). These might include alternative measures, including duty of care approaches (discussed below) and digital well-being indicators (OECD 2021b).

## **10 ESTABLISH COMPREHENSIVE REGULATORY FRAMEWORKS FOR EDUCATIONAL AI**

*I think if this technology goes wrong, it can go quite wrong. And we want to be vocal about that. We want to work with the government to prevent that from happening.*

— Sam Altman, Chief Executive Officer, OpenAI, 2023

As AI becomes integrated into education systems, education experts, teachers, and at times students themselves advocate for regulatory frameworks ensuring that AI is implemented equitably, responsibly, ethically, and in educationally valid ways. Such frameworks would align AI governance with existing legal protections for student rights and privacy, allowing innovations to advance learning without undermining fundamental safeguards.

Countries have different governmental structures and legal systems, requiring flexible regulatory approaches. The following options offer both different examples of regulation as well as distinct implementation mechanisms that can be adapted to various contexts.

| To establish comprehensive regulatory frameworks for educational AI, **adopt a whole-of-government approach to educational AI.**

Governments can embrace a whole-of-government approach that coordinates AI policy across sectors, aligns multiple legislative frameworks, balances competing interests like innovation and regulation, and manages complex cross-sectoral issues (Miao and Holmes 2023).

This comprehensive strategy can involve establishing a national body to lead cross-sectoral coordination and align frameworks with existing legislation on data protection, Internet security, and citizen data security while assessing whether current regulations require adaptation for AI-specific issues. It balances regulation with innovation by promoting cooperation among companies, educational institutions, and public agencies to develop trustworthy models, encouraging open-source ecosystems for sharing computing resources and datasets, and fostering practical AI applications for public good.

This approach also establishes principles for assessing and categorizing AI efficacy, safety, and security before deployment and throughout system life cycles, using risk-based classification mechanisms similar to the European Union Artificial Intelligence Act (EU AI Act) that range from strict regulations banning unacceptable-risk applications to general regulations for lower-risk systems (EU Artificial Intelligence Act 2024). It mandates laws protecting users' personal information and combating unlawful data storage, profiling, and sharing, while enforcing age limits for AI applications primarily designed for adults, which pose substantial risks to children including exposure to inappropriate content and manipulation (Miao and Holmes 2023, 20–21).

| To establish comprehensive regulatory frameworks for educational AI, **require regulation by design.**

Governments can require technology products to meet safe and ethical design standards through regulation by design (RbD). RbD employs a more narrow focus than the whole-of-government approach mentioned above. It integrates regulatory objectives directly into technical design specifications rather than applying them through compliance checks (Mahari and Pataranutaporn 2025, 3). AI designers and developers comply by translating legal requirements into technical elements embedded in system code, pursuing goals that include promoting ethical AI use, student well-being, effective governance, accuracy, transparency, legal protections, autonomy, justice, safety, security, and explainability for developers, users, and society as a whole.

The European Union has pioneered RbD through the GDPR, resulting in technology companies doing business in the EU designing products with stronger embedded privacy protections. While RbD faces challenges including reduced individual autonomy, "consent fatigue," and at times a lack of flexibility, for student-facing AI applications, RbD can transform AI systems into inherently safer platforms by shifting responsibility for safety directly to developers rather than to users (in the case of education, educators, students, or parents) (Prifti et al. 2024). Embedding protections reduces reliance on after-the-fact enforcement mechanisms while transforming safety compliance from an external constraint into an intrinsic feature that can facilitate trust without impinging on innovation (Mahari and Pentland 2024).

| To establish comprehensive regulatory frameworks for educational AI, **enact comprehensive AI-specific legislation.**

Governments can develop legislative frameworks that directly address AI to establish clear requirements and accountability mechanisms that are distinct from broader regulatory coordination or design-level interventions. The EU AI Act, mentioned previously and passed in 2024, provides a model of AI-specific legislation countries can adopt. The Act sorts AI applications into four risk levels—minimal, limited, high,

and unacceptable—with corresponding regulatory requirements. Under the EU AI Act, eight practices threatening safety and rights are prohibited, including harmful manipulation, social scoring, and unauthorized biometric recognition; furthermore, the act mandates conformity assessments for high-risk systems and establishes oversight mechanisms to enforce compliance (European Commission 2025). Such purpose-built legislation sets binding legal requirements from the outset instead of adapting existing laws designed for pre-AI technologies.

| To establish comprehensive regulatory frameworks for educational AI, **establish technical standards and independent auditing requirements.**

Rather than relying solely on prescriptive legislation or design mandates, governments can develop minimum safety standards against which AI systems used in education are independently audited. In the U.S., the National Institute of Standards and Technology, a federal agency within the Department of Commerce that develops measurement, standards, and technology, could establish and oversee such audits.

To ensure compliance with auditing standards, the federal government could employ various mechanisms ranging from restrictive to incentive-based approaches, including prohibiting federal funding recipients from purchasing unaudited AI systems, requiring audits as a condition for federal funding, providing additional funding for districts that purchase audited systems, creating safe harbor legal protections for compliant educational AI tools, withholding funds from school districts experiencing problems with unaudited systems, or funding post-deployment audits (DiPaola et al. 2024, 5).

Each approach involves different trade-offs between enforcement of certainty and flexibility, but at a minimum, education authorities can ensure that schools and the regional and district jurisdictions in which they operate have access to audit standards, results, and information about their educational implications to support informed procurement decisions.

| To establish comprehensive regulatory frameworks for educational AI, **employ “duty of care” laws.**

A “duty of care” represents a legal and ethical obligation requiring individuals or organizations such as schools to take reasonable steps to prevent foreseeable harm to students’ safety or well-being. Governments can impose duty of care laws on technology companies so that student-facing products are designed based on what is best for minors (defined as anyone under the age of 18) as opposed to what is best for the technology company (Bernstein 2025). When governments impose duty of care laws on technology companies, these mandates establish liability to ensure safety and exercise care in design implementation. For minors, this obligation specifically addresses preventing harms, including compulsive use, anxiety, depression, eating disorders, and predictable emotional harm (Bernstein 2023, 2025).

The UK pioneered this regulatory approach in 2020, asserting that technology companies had moral and legal responsibility in terms of how their products affected minors (Department for Science, Innovation and Technology et al. 2022; Haidt 2024). One example of an action incorporated under a duty of care framework involves technology companies setting privacy protections set to the highest default standards (Haidt 2024, 232–233). When applied to social media platforms, duty of care laws typically restrict data collection and use practices, mandate these default privacy protections, and prohibit algorithms that push harmful content to children.

Duty of care frameworks are now expanding beyond social media and could be extended to AI companion platforms (Bernstein 2025). Several U.S. states have already begun such implementation. Vermont prohibits using personal data or designing products that result in predictable emotional harm or compulsive use (State of Vermont 2025). North Carolina goes further by imposing a duty of loyalty that specifically prevents AI bots from creating emotional dependence (Bernstein 2025).

To establish comprehensive regulatory frameworks for educational AI, **update national data sovereignty laws to protect user data.**

AI companies have exploited weak data protection laws in low- and middle-income countries, extracting valuable data resources while building trillion-dollar valuations (Dosunmu 2025). Concerned about one-way flows of economic benefits to “hyperscalers”—companies that provide large-scale cloud computing, networking, and data store—Global South governments can follow the examples of India, Nigeria, Vietnam, and South Africa and update their data sovereignty laws to require local data storage and greater control before permitting AI deployments or data center construction (Dosunmu 2025).

Data sovereignty enables nations to protect their citizens’ privacy, ensure security, and foster local digital economies by controlling data generated within their borders. Unlike traditional AI solutions hosted on hyperscaler platforms under foreign jurisdictions like the U.S. Clarifying Lawful Overseas Use of Data (CLOUD) Act, sovereign AI keeps data within national borders for both storage and processing, eliminating exposure risks through shared services (Broadcom 2025). This approach facilitates compliance with frameworks like GDPR while enabling governments and education systems to select solutions matching their specific requirements (Letort and Linask-Goode 2025).

Miao and Holmes (2023) advocate for stronger national data ownership frameworks that regulate AI providers and prevent exclusive exploitation by tech companies, ensuring mutual benefit from citizen-generated datasets used commercially. Supporting these data sovereignty principles, the Frequency platform offers an open-source infrastructure built on the Decentralized Social Networking Protocol, an open-source standard, which like HTML, is application agnostic. This platform enables all users to control their data through portable digital identities, prioritizing consent and eliminating platform lock-in (Frequency 2025; DSNP.org. 2025).

## **11** **PROCURE TECHNOLOGY THAT PROTECTS STUDENTS’ PRIVACY, SAFETY, AND SECURITY.**

Large education systems wield significant influence over technology companies’ privacy practices through their substantial purchasing power. They can leverage this influence for student privacy, safety, and security by only purchasing platforms, tools, and systems that conform to data protection and privacy principles.

To protect privacy, safety, and security, **use child-friendly procurement criteria.**

Schools can use child-friendly criteria, like those described in Box 30, to make procurement decisions. Students can be empowered to support this procurement process. For example, within a school, a group of students might be trained to read vendor privacy summaries and flag concerns in plain language. These concerns can be assessed as the school makes final procurement decisions, ensuring student involvement in procurement due diligence.

## Child-friendly procurement criteria

A growing number of education organizations and safety experts, most notably the Learning Accelerator (2025), recommend that schools only purchase and use technology products that conform to the following guidelines:

**Technology products that apply *privacy by default*.** Privacy by default ensures that the strictest privacy settings are automatically applied to users by default, requiring them to take action to reduce their privacy protection (Newman 2021).

**Chatbot platforms that employ safety features.** These controls direct AI on how to respond to a user and “receive notifications when the system detects their teen is in a moment of acute distress.” While far from perfect, such controls are an important step in ensuring student safety on online platforms (Hill 2025).

**Tools that prevent data scraping and that charge a fee to or block AI systems from accessing school websites.** Strict access controls and pricing can discourage or block AI “crawlers” from accessing student portals, learning management systems, or platforms with student-generated content. This gives institutions greater control over protecting student information from AI training datasets and alerts them when AI systems attempt to access student content (Allen and Newton 2025).

**Tools with solid data encryption and security standards** that ensure that only data necessary for the function of the platform are collected (Learning Accelerator 2025).

**Tools with transparent data policies** that outline how users’ personal data are collected, stored, shared, and deleted. This can include clear plans for how to address data breaches, ensure that data are kept no longer than needed for its intended purpose, and permanently delete student data when they are no longer needed (Learning Accelerator 2025).

**Tools that allow schools to keep ownership** of student or teacher data and to delete that data when appropriate (Learning Accelerator 2025).

**Tools that continuously update privacy policies** to identify areas for improvement and adjustments and to guide procurement. This can allow education systems to verify compliance with privacy and security standards and ensure that educational technology vendors maintain data privacy practices (Learning Accelerator 2025).

**Guardrails that protect student data.** Education systems can access open-source libraries that allow schools to easily add guardrails or technical controls to AI systems to protect student privacy, safety, and well-being. For example, Hugging Face’s Chatbot Guardrails Arena stress-tests LLMs and privacy guardrails to prevent sensitive data leaks. Nvidia built NeMo Guardrails, an open-source toolkit for adding programmable guardrails to LLM applications. Guardrails AI offers similar open-source functionality. LangChain provides a guardrails library on Github that helps organizations quickly integrate guardrails into operations (McKinsey & Company 2024).

To protect privacy, safety, and security, **incentivize technology companies through certification to design products supporting the learning sciences and children’s well-being.**

Education systems, governments, and families could all use their purchasing power to make child-friendly certification of AI products meaningful. To do this, consumers need clear signals to identify which products are optimized for children and meet child-safety standards. This can be accomplished by certifying products as well as certifying product developers.

**BOX 31**

## Product certifications

**Product certification and criteria.** Anyone purchasing AI products and tools for education can use some of the existing product certifications often run by nonprofits to guide their purchases. Consumers (families, education systems, governments) can consult certifications (often run by nonprofits) and criteria for AI products to help guide their choices. The table below includes a few examples of certifications, frameworks, guidelines, and criteria that can help schools purchase AI products that promote students’ safety and well-being.

ORGANIZATION	FRAMEWORK NAME	SAMPLE CRITERIA
Digital Promise ( <i>Digital Promise 2025</i> )	Responsibly Designed AI (product certification)	<ul style="list-style-type: none"> <li>• Public privacy policy</li> <li>• Public data security policy</li> <li>• Documentation of bias mitigation processes</li> <li>• AI-generated content is labeled</li> <li>• Allows educators/users to override AI decisions</li> </ul>
National Education Association ( <i>NEA 2025</i> )	Vetting AI Resources: A Guide for Educators	<ul style="list-style-type: none"> <li>• Human-centered approach</li> <li>• Evidence-based effectiveness</li> <li>• Ethical and transparent practices</li> <li>• Accessibility and equity</li> <li>• Professional development and support</li> <li>• Privacy, security, and accountability</li> </ul>
Imagine Learning ( <i>Imagine Learning 2025</i> )	Essential AI Tool Vetting Checklist for Educators	<ul style="list-style-type: none"> <li>• Serves a legitimate instructional purpose</li> <li>• Handles student data and privacy transparently and in compliance with regulations</li> <li>• Is equitable and free from harmful bias</li> <li>• Is accessible to all students</li> <li>• Has support and oversight processes in place</li> </ul>
Southern Regional Education Board ( <i>Southern Regional Education Board 2025</i> )	AI Tool Procurement, Implementation, and Evaluation Checklist	<ul style="list-style-type: none"> <li>• Aligns with instructional goals and complements existing curriculum</li> <li>• Intentionally mitigates biases</li> <li>• Complies with best practices in security, ethics, and data use</li> <li>• Allows human oversight</li> <li>• Proven to be effective</li> <li>• Is easily distributed and used</li> </ul>

Researchers and foundations can help support new certifications where gaps exist. For example, the non-profit Common Sense Media's assessment of AI tools could be expanded to serve more families and schools (Common Sense Media, n.d.).

**Developer certification.** Certifying products can be challenging when technology changes so rapidly. But one way to improve the design of AI tools is to help the developers better understand children's learning. As one engineer in a technology company shared with researchers, "We know how technology works, but we don't know how kids work." Technology companies that make products used by children can require that their developers and designers be trained on child development, the science of learning, and essential educational principles. Appropriate training can be developed in partnership with researchers, universities, and governments; where relevant training exists, it can be adapted for different use cases (LearnLab, n.d.).

| To protect privacy, safety, and security, **adopt a no-regret approach.**

"No regret" actions are strategic proactive decisions that yield risk management benefits regardless of how future scenarios unfold. Education systems can borrow this approach from other fields that have successfully employed it and implement pre-emptive measures to protect their technology and AI systems. The potential upside of "no regret" approaches is significant while their downside is minimal because they provide positive outcomes even when assumptions and fears do not materialize (Chorev and Predd 2025). One example of a "no regret" action is a comprehensive audit of every technology vendor and digital tool in an education system. Such action would reveal precisely what student data is collected, how it is used, stored, and accessed. Such an audit would identify potential safety vulnerabilities, enable informed decisionmaking, position schools for regulatory compliance, and respond more quickly to data breaches or cyberattacks because they are responding from a position of knowledge versus ignorance.

## 12 SUPPORT FAMILIES TO MANAGE CHILDREN'S AI USE AT HOME

Families, civil society organizations, and children themselves have an important role to play in helping children stay safe and healthy with their home technology use. At home, parents and caregivers face the double burden of having to help their children learn and grow in both the physical and virtual worlds. Many technology companies with student-facing products outsource a large share of the responsibility for keeping children and youth safe to busy families. Today, some of the most prominent ways in which caregivers are involved in the discussion on AI and children's well-being is from parents whose children have been victims of online harm and abuse (Chatterjee 2025). In addition to technology companies making safe products and governments enforcing safety standards, as discussed above, others can play a role.

| To manage children's AI use at home, **support students to inform families about AI use in schools.**

Students themselves can communicate this information. For example, students may host quarterly forums with the school community that explain what tools are in use, what data flows where, and how to opt out of data sharing where possible. Education systems can prioritize this transparent sharing, including sending basic information home and integrating it into regular family-facing communication.

To manage children’s AI use at home, **provide families and children with safety information.**

Civil society organizations and associations that interface with and support families and children can all play a role in sharing information on how to protect children from online harms. Groups such as parent networks, student groups, pediatrician associations, religious leaders, and sports teams with coaches all have a role to play. There are a wide range of family- and student-facing materials that include things like protecting children’s data, spotting deep fakes, avoiding cyber bullying and sexual exploitation, treating gaming and shopping addiction, limiting exposure to violence and graphic content, and being aware of extortion tactics with AI (see Box 32).

## BOX 32

# Family-facing AI safety information

## Examples of AI safety material for parents

### Parents’ Ultimate Guide to Generative AI

Common Sense Media explains how AI works, how teenagers are using it, and its key risks and benefits; the guide offers parents practical advice for talking with children and youth and guiding safe use (Common Sense Media 2024).

### The Safe AI for Children Alliance: A Comprehensive Guide to AI Risks to Children

The Safe AI for Children Alliance explains the various risks that AI poses to children—from exposure to harmful content, data exploitation, and online grooming to psychological effects, algorithmic bias, and long-term existential threats—and presents a call for urgent action by parents, educators, policy-makers, and society to protect and prepare the next generation (Safe AI for Children 2024).

### A Guide to Cyber Safety for Kids at Every Age

The Singapore government provides parents with clear explanations and checklists of how to help keep children safe in an online world, including privacy settings, discussions to have with children, and behavior shifts parents themselves can make (Infocomm Media Development Authority 2025).

### Artificial Intelligence for Children: Toolkit

The World Economic Forum informs parents and guardians about key factors to consider when purchasing AI-powered products to help them make safe and informed choices (World Economic Forum 2022).

### AI and Education: A Guide for Parents

HP México explains how AI can enhance children’s learning through personalized interactive tools while emphasizing the crucial role of parents in balancing technology use with supervision and age-appropriate guidance (HP 2025).

To manage children’s AI use at home, **limit screen time, including AI use.**

Families and students can talk about screen time and the trade-offs that come with extensive screen time use such as in-person connection, exercise, and sleep. They can engage children in conversations to discuss why AI friends and LLMs are not actually children's companions like their in-person friends. Together, caregivers and children can develop a plan for limiting screen time at home and review the content of children's online interactions. And because parents and caregivers are children's primary educator in life, parents and caregivers can model healthy, responsible technology use.

To manage children's AI use at home, **students can talk with peers.**

Students can talk with their peers about what seem to be good and bad uses of AI. They can support each other in balancing real and online world time and interactions, as well as join student-led organizations and networks that focus on how to lead healthy and safe lives in an online world. Students can join existing youth advocacy groups or start their own if they cannot find one in their community.

### BOX 33

## Student led, peer-to-peer AI discussion

Student Groups Addressing AI Safety

### Design It For Us 2025 AI Policy Platform

Design It For Us is a youth-led organization that aims to ensure that artificial intelligence systems are developed and governed with safety, transparency, meaningful youth participation, educational preparedness, and human-centered labor policies so that young people's rights and well-being are protected in the AI era (Design It For Us 2025).

### Youth in AI

Youth in AI is a youth-led nonprofit based in Africa that empowers young people to become creators, researchers, and ethical leaders in AI by offering education, mentorship, innovation labs, and community building across more than 15 African countries (Youth in AI 2025).

## Conclusion: Navigating the crossroads

The arc of AI is far from complete. As a relatively new technology, the speed and scale of its reach and effects are disorienting, hard to predict, and difficult to fully recognize—and we are all still learning about AI. None of us have all the answers right now.

And yet the stakes to get AI right for our children could not be higher.

As this report has emphasized we can already recognize AI's potential benefits on children's learning, but we also see its risks. Without deliberate interventions by all of us—technology companies, governments, educators, families, researchers, education leaders, and students—the potential harms of AI in education risk undercutting its benefits.

These risks emerge from human choices rather than technological inevitability. They stem from failing to view equity as requisite for quality, lacking guidelines for AI in education, excluding educators from product development, and prioritizing engagement over student well-being. Without immediate intervention, we risk creating students who trade deep understanding for surface efficiency, critical thinking for algorithmic dependence, human companionship for AI relationships, and creative problem-solving for rapid task completion.

Our three-pillar framework—Prosper, Prepare, and Protect—ensures that AI promotes educational excellence, enriches rather than diminishes student learning, and maintains the human-centered focus critical to meaningful education. While discrete, these pillars are essential and interconnected, depending on collective action. As we've argued, education systems can implement thoughtful AI rollout plans aligned with educational objectives. Technology companies can develop ethical tools that support human relationships. Governments can establish flexible guidelines safeguarding student welfare and academic integrity. Families, educators, students, and civil society can model productive AI use while empowering active creation over passive consumption.

We are living in the age of AI. This is the world our children and their children will inhabit, and it will likely impact every facet of their lives. Through coordinated action and unwavering focus on children's development, as we have argued here, we can ensure that this world remains one centered on the needs and aspirations of human beings not algorithms – amplifying rather than replacing human potential, serving rather than substituting for understanding, and preparing every generation to prosper and thrive with these powerful generative tools.

# ENDNOTES

- 1 For example, rule-based systems follow predetermined logic, machine learning models identify patterns from data, and generative AI creates novel content through complex neural networks trained on large data sets (Bergmann 2025).
- 2 Throughout the remainder of this report, the data presented (typically as percentages) refers to the proportion of participants within each group who mentioned each theme. To ensure equal weight was given to each person's voice regardless of how frequently they raised particular concerns, percentages were calculated by dividing the number of participants who mentioned each theme (numerator) by the total number of participants in that group (denominator). This provides a measure of consensus within each stakeholder group rather than simply the volume of comments.
- 3 A/B testing is a controlled randomized test that compares two versions of a digital product or service to determine which performs better. It involves splitting a live audience into two groups, with one group (the control) using version A and the other (the treatment) using version B. By measuring key metrics, technology companies can use data to make decisions about design changes, offers, or user experiences.
- 4 There does seem to be some flexibility in this principle that makes it more compatible with AI practices. For example, AI companies may be allowed to reuse personal data for purposes aligned with those originally communicated or for statistical purposes, assuming there are no "unacceptable risks for the data subject" (Sartor and Lagioia 2020).
- 5 The percent of the population to speak each language was calculated by dividing Ethnologue's 2025 data on the number of speakers (native or second language) of each language by the United Nations Population Divisions' 2025 estimate for total world population.

# REFERENCES

- 1EdTech. 2025. "About 1EdTech." 1EdTech, October 26. <https://www.1edtech.org/about/1edtech>.
- 5Rights Foundation. 2025. "Children and AI Design Code: A Protocol for the Development and Use of AI Systems That Impact Children." 5Rights Foundation. [https://5rightsfoundation.com/wp-content/uploads/2025/03/5rights\\_AI\\_CODE\\_DIGITAL.pdf](https://5rightsfoundation.com/wp-content/uploads/2025/03/5rights_AI_CODE_DIGITAL.pdf).
- Abbas, Muhammad, Farooq Ahmed Jam, and Tariq Iqbal Khan. 2024. "Is It Harmful or Helpful? Examining the Causes and Consequences of Generative AI Usage among University Students." *International Journal of Educational Technology in Higher Education* 21 (10). <https://doi.org/10.1186/s41239-024-00444-7>.
- Abdul Latif Jameel Poverty Action Lab (J-PAL). 2024. "Artificial Intelligence to Strengthen High School Students' Writing Skills." <https://www.povertyactionlab.org/case-study/artificial-intelligence-strengthen-high-school-students-writing-skills>.
- Acar, Oguz A., Phyliss Jia Gai, Yanping Tu, and Jiayi Hou. 2025. "Research: The Hidden Penalty of Using AI at Work." *Harvard Business Review*, August 1. <https://hbr.org/2025/08/research-the-hidden-penalty-of-using-ai-at-work>.
- Adam, David. 2025. "Supportive? Addictive? Abusive? How AI Companions Affect Our Mental Health." *Nature*, May 6. <https://www.nature.com/articles/d41586-025-01349-9>.
- Adams, Kirkwood, and Maria G Baker. 2025. "Characterizing ChatGPT's Feedback for FYW: Analyzing Feedback Responses to Inquiry-Driven Essays." *Thresholds in Education* 48 (2): 159–181.
- Ademokun, Oluwatosin Victoria. 2025. "AI and the GDPR: Understanding the Foundations of Compliance." *TechGDPR* (blog), June 4. <https://techgdpr.com/blog/ai-and-the-gdpr-understanding-the-foundations-of-compliance/>.
- Ahuja, Manoj, Joseph Emmanuel, Biswajit Saha, et al. n.d. "Artificial Intelligence Integration for School Curriculum." CBSE Academic. Accessed December 2, 2025. [https://cbseacademic.nic.in/web\\_material/Curriculum20/AI\\_Integration\\_Manual\\_Introduction.pdf](https://cbseacademic.nic.in/web_material/Curriculum20/AI_Integration_Manual_Introduction.pdf).
- Akiba, Motoko, Gerald K. LeTendre, and Jay P. Scribner. 2007. "Teacher Quality, Opportunity Gap, and National Achievement in 46 Countries." *Educational Researcher* 36 (7): 369–87. <https://doi.org/10.3102/0013189X07308739>.
- Alazemi, Asmaa Falah Theiyab. 2024. "Formative Assessment in Artificial Integrated Instruction: Delving into the Effects on Reading Comprehension Progress, Online Academic Enjoyment, Personal Best Goals, and Academic Mindfulness." *Language Testing in Asia* 14 (44). <https://doi.org/10.1186/s40468-024-00319-8>.
- Alexander, Robin. 2018. "Developing Dialogic Teaching: Genesis, Process, Trial." *Research Papers in Education* 33 (5): 561–98. <https://doi.org/10.1080/02671522.2018.1481140>.
- Ali, Fayaz, Qingyu Zhang, Muhammad Zubair Tauni, and Khuram Shahzad. 2024. "Social Chatbot: My Friend in My Distress." *International Journal of Human-Computer Interaction* 40 (7): 1702–12. <https://doi.org/10.1080/10447318.2022.2150745>.
- Ali, Jamal Kaid Mohammed, Muayad Abdulhalim Ahmad Shamsan, Taha Ahmed Hezam, and Ahmed Majeed. 2023. "Impact of ChatGPT on Learning Motivation: Teachers and Students' Voices." *Journal of English Studies in Arabia Felix* 2 (1): 41–49. <https://doi.org/10.56540/jesaf.v2i1.51>.

- Allen, Will, and Simon Newton. 2025. "Introducing Pay per Crawl: Enabling Content Owners to Charge AI Crawlers for Access." *The Cloudflare Blog*, July 1. <https://blog.cloudflare.com/introducing-pay-per-crawl/>.
- American Psychological Association. 2018. "APA Dictionary of Psychology." American Psychological Association. <https://dictionary.apa.org/cognitive-development>.
- American Psychological Association. 2025a. "Artificial Intelligence and Adolescent Well-Being." American Psychological Association. <https://www.apa.org/topics/artificial-intelligence-machine-learning/health-advisory-ai-adolescent-well-being>.
- American Psychological Association. 2025b. "Misinformation and Disinformation." American Psychological Association. <https://www.apa.org/topics/journalism-facts/misinformation-disinformation>.
- Anderson, Jenny. 2025. "Parents Are Torn about AI in Schools. Here's a Tool That Might Really Help Kids Learn." *How to Be Brave* (blog), November 26. <https://howtobebrave.substack.com/p/parents-are-torn-about-ai-in-schools>.
- Anderson, Jenny, and Rebecca Winthrop. 2025. *The Disengaged Teen: Helping Kids Learn Better, Feel Better, and Live Better*. Penguin Random House.
- Anderson, Monica. 2024. "Americans' View of Technology Companies." Pew Research Center, April 29. <https://www.pewresearch.org/internet/2024/04/29/americans-views-of-technology-companies-2/>.
- Ang, Shermaine. 2024. "Students Are Taught to Use AI Ethically and Responsibly at Different Levels: Chan Chun Sing." *The Straits Times*, January 10. <https://www.straitstimes.com/singapore/politics/students-are-taught-to-use-ai-ethically-and-responsibly-at-different-levels-chan-chun-sing>.
- ANI News. 2022. "'FotonVR' Makes Headway towards Going Global!" ANI News. <https://www.aninews.in/news/business/business/fotonvr-makes-headway-towards-going-global20221124202344/>.
- Anthropic. 2025. "Tracing the Thoughts of a Large Language Model." Anthropic, March 27. <https://www.anthropic.com/research/tracing-thoughts-language-model>.
- Apple, Sam. 2025. "My Couples Retreat With 3 AI Chatbots and the Humans Who Love Them." *Wired*, June 26. <https://www.wired.com/story/couples-retreat-with-3-ai-chatbots-and-humans-who-love-them-replika-nomi-chatgpt/>.
- Arias-Sarah, Pablo, Daniel Bedoya, Christoph Daube, Jean-Julien Aucouturier, Lars Hall, and Petter Johansson. 2024. "Aligning the Smiles of Dating Dyads Causally Increases Attraction." *Proceedings of the National Academy of Sciences* 121 (45): e2400369121. <https://doi.org/10.1073/pnas.2400369121>.
- Arifin, Kasman, Muhammad Sirih, Asmawati Munir, Jahidin Jahidin, and Murni Sabilu. 2024. "The Influence of Multimodal Learning Strategies on Prospective Biology Teachers' Literacy-Numeracy Learning Outcomes." *Eurasia Journal of Mathematics, Science and Technology Education* 21 (1): em2563. <https://doi.org/10.29333/ejmste/15802>.
- Asian College of Teachers. 2025. "China mandates AI literacy: global impacts and lessons for school leaders." Asian College of Teachers, August 1. <https://www.teacherstrainingchina.com/blog/367-China-Mandates-AI-Literacy-Global-Impacts-and-Lessons-for-School-Leaders-blog.php>.
- Australian Government Department of Education. 2025. *China's Approach to AI Education in Schools (Year 1-12)*. <https://www.education.gov.au/download/19495/chinas-approach-ai-education-primary-and-secondary-schools/41945/document/pdf>.

- Australian Government e-Safety Commissioner. 2025. "Safety by Design." eSafety Commissioner. <https://www.esafety.gov.au/industry/safety-by-design>.
- Avelar, Marina, Arushi Terway, Marina Dreux Frotte, and NORRAG. 2020. *Innovative Financing For Education: A Systematic Literature Review*. No. 11. NORRAG. <https://resources.norrageducation.org/resource/595/innovative-financing-for-education-a-systematic-literature-review>.
- Ayers, John W., Adam Poliak, Mark Dredze, et al. 2023. "Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum." *JAMA Internal Medicine* 183 (6): 589–96. <https://doi.org/10.1001/jamainternmed.2023.1838>.
- Bai, Jing, Heqing Huang, and Huahong Ouyang. 2022. "Effects of Group-Play Moderate to Vigorous Intensity Physical Activity Intervention on Executive Function and Motor Skills in 4- to 5-Year-Old Preschoolers: A Pilot Cluster Randomized Controlled Trial." *Frontiers in Psychology* 13 (June). <https://doi.org/10.3389/fpsyg.2022.847785>.
- Bainbridge, Lisanne. 1983. "Ironies of Automation." *Automatica* 19 (6): 775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8).
- Bakir, Vian, and Andrew McStay. 2025. "Move Fast and Break People? Ethics, Companion Apps, and the Case of Character.ai." Social Science Research Network Scholarly Paper No. 5159928. February 28. <https://doi.org/10.2139/ssrn.5159928>.
- Bali, Maha. 2024. "Different Critiques of AI in Education." *Reflecting Allowed* (blog), December 22. <https://blog.mahabali.me/educational-technology-2/different-critiques-of-ai-in-education/>.
- Barrett, Alex, and Austin Pack. 2023. "Not Quite Eye to A.I.: Student and Teacher Perspectives on the Use of Generative Artificial Intelligence in the Writing Process." *International Journal of Educational Technology in Higher Education* 20 (November): 59. <https://doi.org/10.1186/s41239-023-00427-0>.
- Bastani, Hamsa, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakçı, and Rei Mariman. 2024. "Generative AI Can Harm Learning." Social Science Research Network Scholarly Paper No. 4895486. July 15. <https://doi.org/10.2139/ssrn.4895486>.
- Bergmann, Dave. 2025. "What Is Machine Learning?" IBM, August 18. <https://www.ibm.com/think/topics/machine-learning>.
- Bernstein, Gaia. 2022. "A Window of Opportunity to Regulate Addictive Technologies." *Wisconsin Law Review* 2022 (Online): 64–84. <https://wlr.law.wisc.edu/a-window-of-opportunity-to-regulate-addictive-technologies/>.
- Bernstein, Gaia. 2023. *Unwired*. Cambridge University Press.
- Bernstein, Gaia. 2025. "We Are Rushing into the Same Mistakes We Made with Social Media." *After Babel* (blog), August 20. <https://www.afterbabel.com/p/dont-repeat-social-media-mistakes>.
- Birhane, Abeba, Vinay Prabhu, Sang Han, and Vishnu Naresh Boddeti. 2023. "On Hate Scaling Laws for Data-Swamps." Preprint, arXiv:2306.13141, June 28, 2023. <https://doi.org/10.48550/arXiv.2306.13141>.
- Blikstein, Paulo, and Marcelo Worsley. 2016. "Multimodal Learning Analytics and Education Data Mining: Using Computational Technologies to Measure Complex Learning Tasks." *Journal of Learning Analytics* 3 (2): 220–38. <https://doi.org/10.18608/jla.2016.32.11>.
- Bloom, Benjamin S. 1984. "The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring." *Educational Researcher* 13 (6): 4–16. <https://doi.org/10.3102/0013189X013006004>.

Bloom, Benjamin S., Max D. Engelhart, Edward J. Furst, Walker H. Hill, and David R. Krathwohl. 1956. *Taxonomy of Educational Objectives: The Classification of Educational Goals*. Longmans, Green and Co. Ltd.

Bodrova, Elena, and Deborah Leong. 2007. *Tools of the Mind: The Vygotskian Approach to Early Childhood Education*. 2nd ed. Pearson/Merrill Prentice Hall.

Botti-Lodovico, Yolanda. 2025. "Shabana Basij-Rasikh and Mati Amin Believe That AI Can Help Restore the Right to Education for Afghan Girls." *Patrick J. McGovern Foundation* (blog), July 15. <https://medium.com/patrick-j-mcgovern-foundation/shabana-basij-rasikh-and-mati-amin-believe-that-ai-can-help-restore-the-right-to-education-for-75ef744ca308>.

Boulay, Benedict du. 2016. "Artificial Intelligence as an Effective Classroom Assistant." *IEEE Intelligent Systems* 31 (6): 76–81. <https://doi.org/10.1109/MIS.2016.93>.

Bozkurt, Aras, Junhong Xiao, Robert Farrow, et al. 2024. "The Manifesto for Teaching and Learning in a Time of Generative AI: A Critical Collective Stance to Better Navigate the Future." *Open Praxis* 16 (4): 487–513. <https://doi.org/10.55982/open-praxis.16.4.777>.

Bracewell. 2025. "Powering Africa's Digital Future: The Challenge of Energy for Data Center Development." Bracewell LLP, April 16. <https://www.bracewell.com/resources/powering-africas-digital-future-the-challenge-of-energy-for-data-center-development/>.

British Dyslexia Association. 2025. "Dyslexia Factsheet." British Dyslexia Association, September 10. <https://www.bdadyslexia.org.uk/resources>.

Broadcom. 2025. "The Future of AI Is Sovereign: Why Data Sovereignty Is the Key to AI Innovation." *Broadcom News and Stories*, July 28. [https://news.broadcom.com/sovereign-cloud/the-future-of-](https://news.broadcom.com/sovereign-cloud/the-future-of-ai-is-sovereign-why-data-sovereignty-is-the-key-to-ai-innovation)

[ai-is-sovereign-why-data-sovereignty-is-the-key-to-ai-innovation](https://news.broadcom.com/sovereign-cloud/the-future-of-ai-is-sovereign-why-data-sovereignty-is-the-key-to-ai-innovation).

Bronfenbrenner, Urie. 1994. "Ecological Modes of Human Development." Vol. 3 of *International Encyclopedia of Education*. 2nd ed. Elsevier. <https://www.ncj.nl/wp-content/uploads/media-import/docs/6a45c1a4-82ad-4f69-957e-1c76966678e2.pdf>.

Bryan, Claire, and Sharon Lurye. 2025. "Schools Use AI to Monitor Kids, Hoping to Prevent Violence. Our Investigation Found Security Risks." *AP News*, March 12. <https://apnews.com/article/ai-school-chromebook-gaggle-goguardian-security-25a3946727397951fd42324139aaf70f>.

Bryk, Anthony S., and Barbara Schneider. 2002. *Trust in Schools: A Core Resource for Improvement*. Russell Sage Foundation. <https://www.jstor.org/stable/10.7758/9781610440967>.

Buçinca, Zana, Maja Barbara Malaya, and Krzysztof Z. Gajos. 2021. "To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-Assisted Decision-Making." *Proceedings of the ACM on Human-Computer Interaction* 5 (CSCW1): 1–21. <https://doi.org/10.1145/3449287>.

Buckley, Jack, Mario Piacentini, and Alina von Davier. 2021. "Towards a New Generation of Assessment." Presented at the OECD Centre for Educational Research and Innovation conference, Digital education for a strong recovery: A forward look, June 9. <https://www.oecd-events.org/digital-education/session/6f-dc3bb9-79ae-eb11-94b3-501ac5921410/towards-a-new-generation-of-assessment>.

Budiyono, Herman, Marzuki, Wiwek Pudjaningsih, Bambang Prastio, and Ahsani Maulidina. 2025. "Exploring Long-Term Impact of AI Writing Tools on Independent Writing Skills: A Case Study of Indonesian Language Education Students." *International Journal of Information and Education Technology* 15 (5): 1003–13. <https://doi.org/10.18178/ijiet.2025.15.5.2306>.

- Burgess, Matt. 2025. "A Single Poisoned Document Could Leak 'Secret' Data Via ChatGPT." *Wired*, August 6. <https://www.wired.com/story/poisoned-document-could-leak-secret-data-chatgpt/>.
- Burgess, Matt, and Natasha Bernal. 2025. "Chatbots Are Pushing Sanctioned Russian Propaganda." *WIRED*, October 27. <https://www.wired.com/story/chatbots-are-pushing-sanctioned-russian-propaganda/>.
- Burns, Mary. 2020. "From Theory to Use: Making Research More Usable and Useful for Educational Practitioners." In *Annual Review of Comparative and International Education 2019*, edited by Alexander W. Wiseman, vol. 39. Emerald Publishing Limited. <https://doi.org/10.1108/S1479-367920200000039007>.
- Burns, Mary. 2021. *Background Paper Prepared for the Global Education Monitoring Report: Technology and Education*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000378951/PDF/378951eng.pdf.multi>.
- Burns, Mary. 2023a. "Barriers and Supports for Technology Integration: Views from Teachers." UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000386070>.
- Burns, Mary. 2023b. *Distance Education for Teacher Training: Modes, Models, and Methods*. 2nd ed. Education Development Center, Inc. <https://www.edc.org/sites/default/files/uploads/EDC-Distance-Education-Teacher-Training.pdf>.
- Burns, Mary. 2024. "Eyes Wide Open: What We Lose from Generative Artificial Intelligence in Education." In *What Education Demands of Us Today or in the Future*. The University of Lisbon. [https://www.researchgate.net/publication/385001882\\_Eyes\\_Wide\\_Open\\_What\\_We\\_Lose\\_from\\_Generative\\_Artificial\\_Intelligence\\_in\\_Education\\_1](https://www.researchgate.net/publication/385001882_Eyes_Wide_Open_What_We_Lose_from_Generative_Artificial_Intelligence_in_Education_1).
- Burns, Mary, and Petra Wiyakti Bodrogini. 2011. "'The Wisdom of Practice': Web 2.0 as a Cognitive and Community-Building Tool in Indonesia." In *Digital Education: Opportunities for Social Collaboration*, edited by Michael Thomas. Palgrave Macmillan US. [https://doi.org/10.1057/9780230118003\\_9](https://doi.org/10.1057/9780230118003_9).
- Burns, Mary, and Mohammad Issack Santally. 2019. *Information and Communications Technologies in Secondary Education in Sub-Saharan Africa: Policies, Practices, Trends and Recommendations*. Mastercard Foundation, November 19. <https://cdn.buttercms.com/vq95NZFyQ7y6a6IAa8zk>.
- Burns, Mary, and Rebecca Winthrop. 2025. "Imagining Failure to Attain Success: The Art and Science of Pre-Mortems." Brookings Institution, August 21. <https://www.brookings.edu/articles/the-art-and-science-of-pre-mortems/>.
- Cai, Alice, Ian Arawjo, and Elena L. Glassman. 2024. "Antagonistic AI." Preprint, arXiv:2402.07350, February 12, 2024. <https://doi.org/10.48550/arXiv.2402.07350>.
- California Consumer Privacy Act of 2018, Title 1.81. [https://leginfo.ca.gov/faces/codes\\_displayText.xhtml?division=3.&part=4.&law-Code=CIV&title=1.81.5](https://leginfo.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&law-Code=CIV&title=1.81.5).
- Caltrider, Jen, Misha Rykov, and Zoë MacDonald. 2024. "Romantic AI Chatbots Don't Have Your Privacy at Heart." Mozilla Foundation, February 14. <https://www.mozillafoundation.org/en/privacynotincluded/articles/happy-valentines-day-romantic-ai-chatbots-dont-have-your-privacy-at-heart/>.
- Cambium Learning Group. 2024. "Cambium Learning Survey Uncovers How K-12 Teachers and Administrators Are Using AI Right Now and How It Is Shaping the Future of Education." Cambium Learning Group, August 6. <https://www.cambiumlearning.com/about-us/news/6062/cambium-learning-survey-uncovers-how-k-12-teachers-and-administrators-are-using-a>.

Capdehourat, Germán, María Eugenia Curi, Victor Koleszar, and Brian Lorenzo. 2024. *Marco referencial para la enseñanza de la inteligencia artificial*. Ceibal. <https://drive.google.com/file/d/1K1y7zkCN-WEcl32lm5056milYDfL4kHes/view>.

Carey, Bridget. 2024. "The First AI-Powered Storytelling Teddy Bear Is Here. I Gave It to My Kids to Test." CNET, July 15. <https://www.cnet.com/tech/services-and-software/the-first-ai-powered-storytelling-teddy-bear-is-here-i-gave-it-to-my-kids-to-test/>.

Carr, David F. 2023. "ChatGPT Is More Famous, but Character.AI Wins on Engagement." Similarweb, March 23. <https://www.similarweb.com/blog/insights/ai-news/character-ai-engagement/>.

Carr, Nicholas. 2010. *The Shallows*. W. W. Norton & Company.

Carr, Nicholas. 2025. *Superbloom: How Technologies of Connection Tear Us Apart*. W. W. Norton & Company.

CAST. n.d. *The Benefits of Universal Design for Learning*. Last accessed December 14, 2025. <https://www.cast.org/wp-content/uploads/2025/04/UDL-Benefits-Evidence-A11y.pdf>.

CBSE (Central Board of Secondary Education). 2024. "Skill Education—Books and Support Material." CBSE. <https://cbseacademic.nic.in/skill-education-books.html>.

Celik, Ismail, Muhterem Dindar, Hanni Muukkonen, and Sanna Järvelä. 2022. "The Promises and Challenges of Artificial Intelligence for Teachers: A Systematic Review of Research." *TechTrends* 66 (March): 616–30. <https://doi.org/10.1007/s11528-022-00715-y>.

Center for Digital Thriving. 2025. "Ten Fresh Insights on Generative AI from Our Teen Advisors." Harvard Graduate School for Education Center for Digital Thriving. [https://digitalthriving.gse.harvard.edu/wp-content/uploads/2025/05/CDT\\_Youth-Advisory-Insights\\_Gen-AI-Memo.pdf](https://digitalthriving.gse.harvard.edu/wp-content/uploads/2025/05/CDT_Youth-Advisory-Insights_Gen-AI-Memo.pdf).

Center for Humane Technology. 2025. "The CHT Perspective." Center for Humane Technology, October 26. <https://www.humanetech.com/the-cht-perspective>.

Center for Internet Security. 2025. *2025 CIS MS-ISAC K-12 Cybersecurity Report: Where Education Meets Community Resilience*. <https://www.cisecurity.org/insights/white-papers/2025-k12-cybersecurity-report>.

Centers for Disease Control and Prevention. 2024. "About Emotional Well-Being." CDC, May 15. <https://www.cdc.gov/emotional-well-being/about/index.html>.

Central Square Foundation. 2025. "AI Samarth." Central Square Foundation, September 2. <https://www.aisamarth.centralsquarefoundation.org/>.

Centre for Evidence and Implementation (CEI). 2025. "Investigating Parent Views on Teen Use of Generative AI." CEI, September 13. <https://www.ceiglobal.org/work-and-insights/investigating-parent-views-teen-use-generative-ai>.

Chango, Wilson, Rebeca Cerezo, Miguel Sanchez-Santillan, Roger Azevedo, and Cristóbal Romero. 2021. "Improving Prediction of Students' Performance in Intelligent Tutoring Systems Using Attribute Selection and Ensembles of Different Multimodal Data Sources." *Journal of Computing in Higher Education* 33 (3): 614–34. <https://doi.org/10.1007/s12528-021-09298-8>.

Character.ai. 2024. "Optimizing AI Inference at Character.AI." Character.ai (blog), June 20. <https://blog.character.ai/optimizing-ai-inference-at-character-ai/>.

Chatterjee, Rihitu. 2025. "Their Teenage Sons Died by Suicide. Now, They Are Sounding an Alarm about AI Chatbots." *NPR*, September 19. <https://www.npr.org/2025/09/19/1061111111>.

Center for Humane Technology. 2025. "The CHT Perspective." Center for Humane Technology, October 26. <https://www.humanetech.com/the-cht-perspective>.

[npr.org/sections/shots-health-news/2025/09/19/nx-s1-5545749/ai-chatbots-safety-openai-meta-characterai-teens-suicide](https://www.npr.org/sections/shots-health-news/2025/09/19/nx-s1-5545749/ai-chatbots-safety-openai-meta-characterai-teens-suicide).

Chaturvedi, Rijul, Sanjeev Verma, Ronnie Das, and Yogesh K. Dwivedi. 2023. "Social Companionship with Artificial Intelligence: Recent Trends and Future Avenues." *Technological Forecasting and Social Change* 193 (August): 122634. <https://doi.org/10.1016/j.techfore.2023.122634>.

Chen, Bodong, Jiayu Cheng, Chen Wang, and Vivian Leung. 2025. "Pedagogical Biases in AI-Powered Educational Tools: The Case of Lesson Plan Generators." *Social Innovations Journal* 30 (2025). <https://socialinnovationsjournal.com/index.php/sij/article/view/10004/8134>.

Child Welfare League Foundation (CWLF). 2025. "CWLF Survey: 58% of Teens Use AI for Schoolwork, but 60% Rarely Verify Accuracy or Safeguard Privacy." Child Welfare League Foundation, September 24. [https://www.children.org.tw/english/news\\_detail/2025-survey-online-safety](https://www.children.org.tw/english/news_detail/2025-survey-online-safety).

Chorev, Matan, and Joel Predd. 2025. "America Should Assume the Worst About AI." *Foreign Affairs*, July 22. <https://www.foreignaffairs.com/united-states/artificial-intelligence-geopolitics-worst-about-ai>.

Chow, Andrew R., and Angela Haupt. 2025. "What Happened When a Doctor Posed as a Teen for AI Therapy." *TIME*, June 12. <https://time.com/7291048/ai-chatbot-therapy-kids/>.

Chowdhury, Sanjana. 2024. "Meet India's Very First AI Humanoid Robot Teacher—Gone Viral." *Asian College of Teachers* (blog), July 10. <https://www.asiancollegeofteachers.com/blogs/1981-Meet-Indias-Very-First-AI-Humanoid-Robot-Teacher---Gone-Viral-blog.php>.

Clun, Rachel. 2025. "Afghan Women Make Do with AI Friends 'to Give Yourself Fake Hope' under Taliban Rule." *The Independent*, March 7.

<https://www.the-independent.com/asia/south-asia/afghanistan-women-ai-taliban-international-b2710878.html>.

Cobing, Laura. 2025. "AI-Powered eLearning Assessments: Adaptive, Gamified, and Data-Driven." *Mindsmith* (blog), March 17. <https://www.mindsmith.ai/blog/ai-powered-elearning-assessments-adaptive-gamified-and-data-driven>.

Code.org. 2025. "Free K–12 Curriculum for Computer Science and AI | Code.Org." October 26. <https://code.org/en-US>.

Code.org, CoSN, Digital Promise, European EdTech Alliance, J. Larimore, and PACE. 2025. *AI Guidance for Schools Toolkit*. TeachAI. <https://www.teachai.org/toolkit>.

Cohen, David K., and Deborah Loewenberg Ball. 1999. *Instruction, Capacity, and Improvement*. Consortium for Policy Research in Education (CPRE) Publications, Philadelphia, PA. <https://eric.ed.gov/?id=ED431749>.

Cohen, Geoffrey. 2022. "Understanding and Overcoming Belonging Uncertainty." *Behavioral Scientist*, October 10. <https://behavioralscientist.org/understanding-and-overcoming-belonging-uncertainty/>.

Collins, Katie. 2025. "Online Age Verification Rules Are Popping Up Everywhere. Here's What You Need to Know." CNET, August 27. <https://www.cnet.com/tech/services-and-software/online-age-verification-rules-are-popping-up-everywhere-heres-what-you-need-to-know/>.

Common Sense Media. 2024. "Parents' Ultimate Guide to Generative AI." Common Sense Media, September 17. <https://www.commonsensemedia.org/articles/parents-ultimate-guide-to-generative-ai>.

Common Sense Media. 2025a. "CSM AI Risk Assessment: Social AI Companions." Common Sense Media, April 10. <https://www.com->

[monsensemedia.org/sites/default/files/pug\\_csm-ai-risk-assessment-social-ai-companions\\_final.pdf](https://monsensemedia.org/sites/default/files/pug_csm-ai-risk-assessment-social-ai-companions_final.pdf).

Common Sense Media. 2025b. "Our AI Initiatives." Common Sense Media, November 21. <https://www.commonsensemedia.org/ai>.

Conner, Jerusha O., and Denise C. Pope. 2013. "Not Just Robo-Students: Why Full Engagement Matters and How Schools Can Promote It." *Journal of Youth and Adolescence* 42 (9): 1426–42. <https://doi.org/10.1007/s10964-013-9948-y>.

Consortium for School Networking (CoSN). 2025. "CoSN Releases 2025 State of EdTech District Leadership Report." CoSN, May 6. <https://www.cosn.org/cosn-news/cosn-releases-2025-state-of-edtech-district-leadership-report/>.

Costello, John M. 2000. "AAC Intervention in the Intensive Care Unit: The Children's Hospital Boston Model." *Augmentative and Alternative Communication* 16 (September): 137–52.

Council of Writing Program Administrators, National Council of Teachers of English, and National Writing Project (CWPA, NCTE, and NWP). 2011. *Framework for Success in Postsecondary Writing*. [https://wpacouncil.org/aws/CWPA/asset\\_manager/get\\_file/350201](https://wpacouncil.org/aws/CWPA/asset_manager/get_file/350201).

Cox Mobile. 2025. *Online Habits and Safety Concerns Across Three Generations*. Cox Mobile. [https://www.coxmobilesafety.com/2025/cox-mobile-online-safety-survey-infographic\\_tips.pdf](https://www.coxmobilesafety.com/2025/cox-mobile-online-safety-survey-infographic_tips.pdf).

Cruz-Jesus, Frederico, Mauro Castelli, Tiago Oliveira, et al. 2020. "Using Artificial Intelligence Methods to Assess Academic Achievement in Public High Schools of a European Union Country." *Heliyon* 6 (6): e04081. <https://doi.org/10.1016/j.heliyon.2020.e04081>.

Cukurova, Mutlu. 2024. "The Interplay of Learning, Analytics and Artificial Intelligence in Education: A Vision for Hybrid Intelligence." *British Journal of*

*Educational Technology* 56 (2): 469–88. <https://doi.org/10.1111/bjet.13514>.

Curriculum Associates. 2025. "AI Labs." Curriculum Associates, October 3. <https://www.curriculumassociates.com/about/ailabs>.

Darling-Hammond, Linda, Lisa Flook, Channa Cook-Harvey, Brigid Barron, and David Osher. 2020. "Implications for Educational Practice of the Science of Learning and Development." *Applied Developmental Science* 24 (2): 97–140. <https://doi.org/10.1080/10888691.2018.1537791>.

Darvishi, Ali, Hassan Khosravi, Shazia Sadiq, Dragan Gašević, and George Siemens. 2024. "Impact of AI Assistance on Student Agency." *Computers & Education* 210 (March): 104967. <https://doi.org/10.1016/j.compedu.2023.104967>.

Day of AI. 2025. "Home." Day of AI, October 6. <https://dayofai.org/>.

Deci, Edward L., and Richard M. Ryan. 2008. "Self-Determination Theory: A Macrotheory of Human Motivation, Development, and Health." *Canadian Psychology/Psychologie Canadienne* 49 (3): 182–85.

Deibert, Ronald, John Palfrey, Rafal Rohozinski, and Jonathan L. Zittrain. 2010. *Access Controlled: The Shaping of Power, Rights, and Rule in Cyberspace*. MIT Press. <https://direct.mit.edu/books/oa-edited-volume/1872/Access-ControlledThe-Shaping-of-Power-Rights-and>.

Dell'Acqua, Fabrizio, Edward McFowland, Ethan R. Mollick, et al. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 24–013, The Wharton School Research Paper. September 27. <https://doi.org/10.2139/ssrn.4573321>.

Delors, Jacques, ed. 1998. *Learning the Treasure Within. Report to UNESCO of the International Commission on Education for the Twenty-First Century*. 2nd ed. UNESCO Publishing.

Denejkina, Anna. 2025. *Young People's Perception and Use of Generative AI 2025*. The Insight Centre and Student Edge. [https://studentedge-content.blob.core.windows.net/documents/youth-insight/2025/05/2025\\_Young\\_Peoples\\_Perception\\_and\\_Use\\_Of\\_Generative\\_AI.pdf](https://studentedge-content.blob.core.windows.net/documents/youth-insight/2025/05/2025_Young_Peoples_Perception_and_Use_Of_Generative_AI.pdf).

Deng, Ruiqi, Maoli Jiang, Xinlu Yu, Yuyan Lu, and Shasha Liu. 2025. "Does ChatGPT Enhance Student Learning? A Systematic Review and Meta-Analysis of Experimental Studies." *Computers & Education* 227 (April): 105224. <https://doi.org/10.1016/j.compedu.2024.105224>.

Department for Education. 2025. "EdChat—The Department's Generative AI Chatbot." Department for Education, South Australia. <https://www.education.sa.gov.au/parents-and-families/curriculum-and-learning/ai/edchat>.

Department for Science, Innovation, and Technology. 2023. "Online Safety Act: Explainer." GOV.UK. <https://www.gov.uk/government/publications/online-safety-act-explainer/online-safety-act-explainer>.

Department for Science, Innovation, and Technology. 2025. "International Scientific Report on the Safety of Advanced AI: Interim Report." GOV.UK, September 15. <https://www.gov.uk/government/publications/international-scientific-report-on-the-safety-of-advanced-ai/international-scientific-report-on-the-safety-of-advanced-ai-interim-report>.

Department for Science, Innovation and Technology, Office for Artificial Intelligence, Department for Digital, Culture, Media & Sport, and Department for Business, Energy & Industrial Strategy. 2022. "Establishing a Pro-Innovation Approach to Regulating AI." GOV.UK, July 20. <https://www.gov.uk/government/publications/>

[establishing-a-pro-innovation-approach-to-regulating-ai](#).

Design It for Us. 2025. "Design It for Us." Design It for Us AI Policy Platform. <https://designitforus.org/platform/design-it-for-us-2025-ai-policy-platform/>.

De Simone, Martin Elias, Federico Hernan Tiberti, Maria Rebeca Barron Rodriguez, Federico Alfredo Manolio, Wuraola Mosuro, and Eliot Jolomi Dikoru. 2025. *From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria*. World Bank Policy Research Working Paper. <https://documents.worldbank.org/en/publication/documents-reports/documentdetail/099548105192529324>.

Deursen, Alexander JAM van, and Jan AGM van Dijk. 2014. "The Digital Divide Shifts to Differences in Usage." *New Media and Society* 16 (3). <https://doi.org/10.1177/1461444813487959>.

Deutscher, R. R., N. C. Holthius, S. I. Maldonado, et al. 2021. *Project-Based Learning Leads to Gains in Science and Other Subjects in Middle School and Benefits All Learners*. Lucas Education Research. <https://www.lucasedresearch.org/publication/project-based-learning-leads-to-gains-in-science-and-other-subjects-in-middle-school-and-benefits-all-learners/>.

Dewey, John. 1916. *Democracy and Education*. Macmillan.

Dhar, Vilas. 2025. "AI's Development Still Depends on All of Us." *Boston Globe*, May 14. <https://www.bostonglobe.com/2025/05/14/opinion/ai-governance-regulation-innovation/>.

Di Mitri, Daniele, Jan Schneider, Marcus Specht, and Hendrik Drachsler. 2018. "From Signals to Knowledge: A Conceptual Model for Multimodal Learning Analytics." *Journal of Computer Assisted Learning* 34 (4): 338–49. <https://doi.org/10.1111/jcal.12288>.

- Diamond, Adele. 2013. "Executive Functions." *Annual Review of Psychology* 64 (January): 135–68. <https://doi.org/10.1146/annurev-psych-113011-143750>.
- Diamond, Adele, and Kathleen Lee. 2011. "Interventions Shown to Aid Executive Function Development in Children 4 to 12 Years Old." *Science* 333 (6045): 959–64. <https://doi.org/10.1126/science.1204529>.
- Dicastery for the Doctrine of the Faith & Dicastery for Culture and Education. 2025. "Antiqua et Nova. Note on the Relationship Between Artificial Intelligence and Human Intelligence." The Vatican, January 28. [https://www.vatican.va/roman\\_curia/congregations/cfaith/documents/rc\\_ddf\\_doc\\_20250128\\_antiqua-et-nova\\_en.html](https://www.vatican.va/roman_curia/congregations/cfaith/documents/rc_ddf_doc_20250128_antiqua-et-nova_en.html).
- Digital Promise. 2025. "Responsibly Designed AI." Digital Promise, October 26. <https://digitalpromise.org/product-certifications/responsibly-designed-ai/>.
- Dijk, Jan van. 2020. *The Digital Divide*. John Wiley & Sons.
- DiPaola, Daniella, Andrés F. Salazar-Gómez, Hal Abelson, Eric Klopfer, David Goldston, and Cynthia Breazeal. 2024. *How Policy Can Help Ensure the Proper Use of AI in K-12 Education*. MIT Media Lab. <https://www.media.mit.edu/publications/how-policy-can-help-ensure-the-proper-use-of-ai-in-k-12-education/>.
- Dixon, Stacy Jo. 2025. "Global Snapchat User Age & Gender Distribution 2025." Statista, May 22. <https://www.statista.com/statistics/933948/snapchat-global-user-age-distribution/>.
- Do, Yeseul. 2025. "Beyond Privacy: Regulating ChatGPT for Young Adults in Educational Contexts." *NYU Journal of Intellectual Property & Entertainment Law* 14 (2): 263–314.
- Doherty, Jonathan. 2020. "A Systematic Review of Literature on Teacher Attrition and School-Related Factors That Affect It." *Teacher Education Advancement Network Journal* 12 (1): 75–84.
- Dosunmu, Damilare. 2025. "Why Big Tech Is Threatened by a Global Push for Data Sovereignty." *Rest of World*, July 9. <https://restofworld.org/2025/big-tech-data-sovereignty/>.
- DSNP.org. 2025. "DSNP—Decentralized Social Networking Protocol." DSNP, October 28. <https://dsn timer.org/>.
- Duff, Lesley. 2018. "Tackling Teacher Workload: Should We Ditch the Data?" National Foundation for Education Research, November 12. <https://www.nfer.ac.uk/blogs/tackling-teacher-workload-should-we-ditch-the-data/>.
- Duffy, Clare. 2024. "'There Are No Guardrails.' This Mom Believes an AI Chatbot Is Responsible for Her Son's Suicide." *CNN*, October 30. <https://www.cnn.com/2024/10/30/tech/teen-suicide-character-ai-lawsuit>.
- Duffy, Clare. 2025. "Parents of 16-Year-Old Adam Raine Sue OpenAI, Claiming ChatGPT Advised on His Suicide." *CNN*, August 27. <https://www.cnn.com/2025/08/26/tech/openai-chatgpt-teen-suicide-lawsuit>.
- Duricic, Anja. 2024. "Small Language Models: A Beginner's Guide." *Ataccama* (blog), November 19. <https://www.ataccama.com/blog/small-language-models>.
- Durlak, Joseph A, Joseph L. Mahoney, and Alaina E. Boyle. 2022. "What We Know, and What We Need to Find Out About Universal, School-Based Social and Emotional Learning Programs for Children and Adolescents: A Review of Meta-Analyses and Directions for Future Research." *Psychological Bulletin* 148 (11-12): 765–82. <https://doi.org/10.1037/bul0000383>.
- Dwyer, Maddy, and Elizabeth Laird. 2024. *Report—Up in the Air: Educators Juggling the Potential of Generative AI with Detection, Discipline, and*

*Distrust*. Center for Democracy and Technology, March 27. <https://cdt.org/insights/report-up-in-the-air-educators-juggling-the-potential-of-generative-ai-with-detection-discipline-and-distrust/>.

e-Estonia. 2025. "Estonia's groundbreaking national initiative: AI Leap programme to bring AI tools to all schools." e-Estonia, February 26. <https://e-estonia.com/estonia-announces-a-groundbreaking-national-initiative-ai-leap-program-to-bring-ai-tools-to-all-schools/>.

Economic Times. 2025. "OpenAI and ARISE Collaborate to Empower K-12 Educators in India with AI Integration." *The Economic Times*, September 14. <https://economictimes.indiatimes.com/tech/startups/openai-and-arise-collaborate-to-empower-k-12-educators-in-india-with-ai-integration/articleshow/123881394.cms>.

Economist. 2024. "Why Don't Women Use Artificial Intelligence?" *The Economist*, August 8. <https://www.economist.com/finance-and-economics/2024/08/21/why-dont-women-use-artificial-intelligence>.

Economist. 2025. "Banning Smartphones in Classrooms Helps Students." *The Economist*, September 4. <https://www.economist.com/united-states/2025/09/04/banning-smartphones-in-classrooms-helps-students>.

Edelman Trust Institute. 2025. "2025 Edelman Trust Barometer: Trust and the Crisis of Grievance with Insights for the Technology Sector." Edelman Trust Institute. [https://www.edelman.com/sites/g/files/aatuss191/files/2025-02/2025%20Edelman%20Trust%20Barometer\\_Insights%20Technology%20Sector\\_FINAL.pdf](https://www.edelman.com/sites/g/files/aatuss191/files/2025-02/2025%20Edelman%20Trust%20Barometer_Insights%20Technology%20Sector_FINAL.pdf).

Education Endowment Foundation (EEF). 2016. "Philosophy for Children." EEF, December 9. <https://educationendowmentfoundation.org.uk/projects-and-evaluation/projects/philosophy-for-children>.

Education Outcomes Fund (EOF). 2025. "The Evidence." EOF, October 26. <https://www.educationoutcomesfund.org/copy-of-backup-the-evidence-old>.

Edwards, Charlotte. 2025. "Meta Investigated over AI Having 'Sensual' Chats with Children." *BBC*, August 18. <https://www.bbc.com/news/articles/c3dpmlvx1k2o>.

Ellipsis. 2025. "High-Demand Radio Frequency Spectrum." Ellipsis, February 26. <https://www.ellipsis.co.za/high-demand-spectrum/>.

Elongué, Christian. 2025. "Major Ways African Languages Are Being Boosted by AI and How LSPs Can Maximize Them." Kabod Group, August 29. <https://kabodgroup.com/ways-african-languages-are-being-boosted-by-ai/>.

Emerson, Ralph Waldo. 1908. *The Essay on Self-Reliance*. The Roycrofters.

Empower Africa. 2025. "Nigeria and Gates Foundation Launch \$7.5 Million AI Scaling Hub to Transform Key Sectors." Empower Africa, June 4. <https://empowerafrica.com/nigeria-and-gates-foundation-launch-7-5-million-ai-scaling-hub-to-transform-key-sectors/>.

Engels, Zoe. 2025. "AI Chat to Support District 112 Teens in Pilot Program." *The Record*, May 1. <https://www.therecordnorthshore.org/2025/04/30/ai-chat-to-support-district-112-teens-in-pilot-program/>.

Epley, Nicholas. 2019. "Humanity Is Carried on the Voice." The University of Chicago Booth School of Business, July 30. <https://www.chicagobooth.edu/review/humanity-carried-voice>.

Epstein, Kate. 2025. "We, Robots." *Persuasion* (blog), January 13. <https://www.persuasion.community/p/we-robots>.

Ethnologue. 2025. "What Are the Top 200 Most Spoken Languages?" *The Ethnologue*. <https://www.ethnologue.com/insights/ethnologue200/>.

EU Artificial Intelligence Act. 2024. "High-Level Summary of the AI Act." Updated February 27. <https://artificialintelligenceact.eu/high-level-summary/>.

European Commission. 2024. "AI Act." European Commission. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>.

European Commission and OECD. 2025. *Empowering Learners for the Age of AI* (review draft). European Commission and OECD. [https://ailiteracyframework.org/wp-content/uploads/2025/05/AI\\_Lit\\_Framework\\_ReviewDraft.pdf](https://ailiteracyframework.org/wp-content/uploads/2025/05/AI_Lit_Framework_ReviewDraft.pdf).

Intersoft Consulting. 2016. "General Data Protection Regulation (GDPR)." Intersoft Consulting. <https://gdpr-info.eu/>.

European Union. 2025. "Council Simulation Game." European Union. [https://learning-corner.learning.europa.eu/learning-materials/council-simulation-game\\_en](https://learning-corner.learning.europa.eu/learning-materials/council-simulation-game_en).

Fan, Yizhou, Luzhen Tang, Huixiao Le, et al. 2024. "Beware of Metacognitive Laziness: Effects of Generative Artificial Intelligence on Learning Motivation, Processes, and Performance." *British Journal of Educational Technology* 56 (2): 489–530. <https://doi.org/10.1111/bjet.13544>.

Fang, Cathy Mengying, Auren R. Liu, Valdemar Danry, et al. 2025. "How AI and Human Behaviors Shape Psychosocial Effects of Chatbot Use: A Longitudinal Randomized Controlled Study." Preprint, arXiv:2503.17473, March 21, 2025. <https://doi.org/10.48550/arXiv.2503.17473>.

Federal Trade Commission. 2013. "Children's Online Privacy Protection Rule ('COPPA')." Federal Trade Commission, July 25. <https://www.ftc.gov/legal-library/browse/rules/childrens-online-privacy-protection-rule-coppa>.

Federal Trade Commission. 2025a. "FTC Launches Inquiry into AI Chatbots Acting as Companions." Federal Trade Commission,

September 11. <https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-launches-inquiry-ai-chatbots-acting-companions>.

Federal Trade Commission. 2025b. "Section 5: Unfair or Deceptive Acts or Practices." Federal Trade Commission. <https://www.federalreserve.gov/boarddocs/supmanual/cch/200806/ftca.pdf>.

Feldstein, Steve. 2019. "The Global Expansion of AI Surveillance." Carnegie Endowment for International Peace, September 17. <https://carnegieendowment.org/research/2019/09/the-global-expansion-of-ai-surveillance>.

Ferman, Bruno, Lycia Lima, and Flávio Riva. 2021. "Artificial Intelligence, Teacher Tasks and Individualized Pedagogy." SocArXiv qw249, February 17. Center for Open Science. <https://ideas.repec.org/p/osf/socarx/qw249.html>.

Findley, Madeleine, Michelle Kallen, and David Layden. 2024. "A Look at 5 States' New Data Privacy Laws." Jenner.com, October 7. <https://www.jenner.com/a/web/2ETrwFPLALdavo2os5Bfz4/law-360-a-look-at-5-states-new-data-privacy-laws.pdf>.

Fischer, Julia Freeland. 2025. "AI Can Make Schools More Human, But Only If Schools Prioritize Relationship Metrics." *Harvard Social Impact Review*, September 12. <https://www.sir.advancedleadership.harvard.edu/articles/ai-can-make-schools-more-human-if-schools-prioritize-relationship-metrics>.

Floridi, Luciano. 2024. "Hypersuasion—On AI's Persuasive Power and How to Deal with It." Center for Digital Ethics (CEDE). .Network Scholarly Paper No. 4815890. May 3. <https://papers.ssrn.com/abstract=4815890>.

fotonVR. 2025. "Virtual Reality for School Education." Foton VR, October 3. <https://fotonvr.com/virtual-reality-for-school-education>.

Foucault, Michel. 1995. *Discipline and Punish: The Birth of the Prison*. 2nd ed. Vintage Books.

Frantz, Erica, Andrea Kendall-Taylor, and Joseph Wright. 2020. "Digital Repression in Autocracies." Varieties of Democracy Institute at the University of Gothenburg and Notre Dame University Working Paper. <https://www.readkong.com/page/fullscreen/digital-repression-in-autocracies-erica-frantz-andrea-9541444>.

Franze, Andrew, Christina R. Galanis, and Daniel L. King. 2023. "Social Chatbot Use (e.g., ChatGPT) among Individuals with Social Deficits: Risks and Opportunities." *Journal of Behavioral Addictions* 12 (4): 871–72. <https://doi.org/10.1556/2006.2023.00057>.

Freire, Paulo. 1970. *Pedagogy of the Oppressed*. Bloomsbury Academic.

Freitas, Julian De, Zeliha Oğuz-Uğuralp, and Ahmet Kaan-Uğuralp. 2025. "Emotional Manipulation by AI Companions." Harvard Business School Working Paper No. 26–005. <http://dx.doi.org/10.2139/ssrn.5390377>.

Frequency. 2025. "Frequency: Empowering Self-Sovereign Digital Identities and User-Controlled Data." Frequency, October 28. <https://www.frequency.xyz/>.

Fukuyama, Francis. 2025. "Delegation and Destruction." *Persuasion* (blog), January 13. <https://www.persuasion.community/p/delegation-and-destruction>.

Fuller, Bruce, and Hoyun Kim. 2022. "Systems Thinking to Transform Schools: Identifying Levers That Lift Educational Quality." Brookings Institution, September 12. <https://www.brookings.edu/articles/systems-thinking-to-transform-schools-identifying-levers-that-lift-educational-quality/>.

Furze, Leon. 2024. "Synthetic Sycophants: Why 'Yes-Bots' Are a Problem for Education." *Leon Furze* (blog), November 10. <https://leonfurze.com/2024/11/11/synthetic-sycophants-why-yes-bots-are-a-problem-for-education/>.

Galinsky, Ellen. 2024. *The Breakthrough Years: A New Scientific Framework for Raising Thriving Teens*. Flatiron Books.

Gallup and Walton Family Foundation. 2025. *Teaching for Tomorrow: Unlocking Six Weeks a Year With AI*. Gallup. [https://nextgeninsights.waltonfamilyfoundation.org/wp-content/uploads/2025/06/Teachers\\_AI\\_Report.pdf](https://nextgeninsights.waltonfamilyfoundation.org/wp-content/uploads/2025/06/Teachers_AI_Report.pdf).

Gallup, Walton Family Foundation, and Heartland Forward. 2025. *Voices of Gen Z: Preparing the Heartland for an AI Future*. [https://heartlandforward.org/wp-content/uploads/2025/05/20250507\\_WFF\\_GenZ\\_HeartlandForward\\_AI\\_Report\\_FINAL.pdf](https://heartlandforward.org/wp-content/uploads/2025/05/20250507_WFF_GenZ_HeartlandForward_AI_Report_FINAL.pdf).

García-Méndez, Silvia, Francisco de Arriba-Pérez, and María del Carmen Somoza-López. 2025. "A Review on the Use of Large Language Models as Virtual Tutors." *Science & Education* 34 (2): 877–92. <https://doi.org/10.1007/s11191-024-00530-2>.

Gardony, Aaron L., Tad T. Brunyé, and Holly A. Taylor. 2015. "Navigational Aids and Spatial Memory Impairment: The Role of Divided Attention." *Spatial Cognition & Computation* 15 (4): 246–84. <https://doi.org/10.1080/13875868.2015.1059432>.

Garfinkle, Adam. 2020. "The Erosion of Deep Literacy." *National Affairs* 65 (Fall 2025). <https://nationalaffairs.com/publications/detail/the-erosion-of-deep-literacy>.

GDPR Advisor. 2023. "GDPR and Artificial Intelligence: Challenges and Ethical Considerations." GDPR Advisor, May 13. <https://www.gdpr-advisor.com/gdpr-and-artificial-intelligence-challenges-and-ethical-considerations/>.

GDPR Advisor. 2024. "AI and Data Privacy: Navigating GDPR in the Age of Machine Learning." GDPR Advisor, August 20. <https://www.gdpr-advisor.com/ai-and-data-privacy-navigating-gdpr-in-the-age-of-machine-learning/>.

Georgetown Law. 2019. "Institute for Public Representation's 2018 Complaint Against YouTube and Google Regarding Protection of Children Results in Settlement, Fine." Georgetown Law, September 5. <https://www.law.georgetown.edu/news/institute-for-public-representations-2018-complaint-against-youtube-and-google-for-online-protection-of-children-results-in-settlement-fine/>.

Gerlich, Michael. 2025. "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking." *Societies* 15 (1): 6. <https://doi.org/10.3390/soc15010006>.

Ghana Natural Language Processing (NLP). 2025. "About Us." Ghana Natural Language Processing. <https://ghananlp.org/about>.

Ghosh, Pallab. 2025. "The People Who Think AI Might Become Conscious." *BBC*, May 26. <https://www.bbc.com/news/articles/c0k3700zljjo>.

Global Partnership for Education. 2023. "Debt2ed." Global Partnership, September. <https://www.globalpartnership.org/node/document/download?file=document/file/2023-09-gpe-fact-sheet-debt2ed-rev.pdf>.

Goldstein, Brett J., and Brett V. Benson. 2025. "The Era of A.I. Propaganda Has Arrived, and America Must Act." *The New York Times*, August 5. <https://www.nytimes.com/2025/08/05/opinion/china-ai-propaganda.html>.

Gorichanaz, Tim. 2023. "Accused: How Students Respond to Allegations of Using ChatGPT on Assessments." *Learning: Research and Practice* 9 (2): 183–96. <https://doi.org/10.1080/23735082.2023.2254787>.

Graham, Steve. 2019. "Changing How Writing Is Taught." *Review of Research in Education* 43 (1): 277–303. <https://doi.org/10.3102/0091732X18821125>.

Grand View Research. 2025. "Education Technology Market (2025–2030)." Grand View Research,

October 3. <https://www.grandviewresearch.com/industry-analysis/education-technology-market>.

Grose, Jessica. 2025. "A.I. Will Destroy Critical Thinking in K-12." *The New York Times*, May 14. <https://www.nytimes.com/2025/05/14/opinion/trump-ai-elementary.html>.

Grosso, David P., Michelle R. Bowling, Starshine S. Chun, and Brooke M. Delaney. 2024. "The Development of AI and Protecting Student Data Privacy." *ArentFox Schiff* (blog), February 21. <https://www.afslaw.com/perspectives/ai-law-blog/the-development-ai-and-protecting-student-data-privacy>.

GSMA. 2016. "Mobile Privacy Principles: Promoting Consumer Privacy in the Mobile Ecosystem." GSM Association. [https://www.gsma.com/solutions-and-impact/connectivity-for-good/public-policy/wp-content/uploads/2012/03/GSMA2016\\_Guidelines\\_Mobile\\_Privacy\\_Principles.pdf](https://www.gsma.com/solutions-and-impact/connectivity-for-good/public-policy/wp-content/uploads/2012/03/GSMA2016_Guidelines_Mobile_Privacy_Principles.pdf).

GSMA. 2020. "Connected Women: The Mobile Gender Gap Report 2020." GSM Association, May. <https://www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-for-development/wp-content/uploads/2020/05/GSMA-The-Mobile-Gender-Gap-Report-2020.pdf>.

GSMA. 2025. "The Mobile Economy 2025." GSM Association, September 5. <https://www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-economy/>.

Guo, Kai, and Deliang Wang. 2024. "To Resist It or to Embrace It? Examining ChatGPT's Potential to Support Teacher Feedback in EFL Writing." *Education and Information Technologies* 29 (7): 8435–63. <https://doi.org/10.1007/s10639-023-12146-0>.

Gustafsson-Wright, Emily, Izzy Boggild-Jones, Dean Segell, and Justice Durland. 2017. "Impact Bonds in Developing Countries: Early Learnings from the Field." Brookings Institution. <https://www.brookings>.

[edu/articles/impact-bonds-in-developing-countries-early-learnings-from-the-field/](#).

Haidt, Jonathan. 2024. *The Anxious Generation: How the Great Rewiring of Childhood Is Causing an Epidemic of Mental Illness*. Penguin.

Hale, Annie, Michael Sears, John Harlow, et al. 2025. "Dreamscape Learn Fall 2022–Spring 2024 Campus Immersion Retrospective." Arizona State University Course Syllabus. <https://drive.google.com/file/d/1usE73Y3hrGf0KeaySKvOUN2rn-79NbkGx/view>.

Han, Ong Hua [MP]. 2024. "Artificial Intelligence." Ministry of Education, Singapore: News, Parliamentary Replies, January 9. <http://www.moe.gov.sg/news/parliamentary-replies/20230109-artificial-intelligence>.

Handa, Kunal, Drew Bent, Alex Tamkin, et al. 2025. "Anthropic Education Report: How University Students Use Claude." Anthropic, April 8. <https://www.anthropic.com/news/anthropic-education-report-how-university-students-use-claude>.

Hao, Karen. 2025. *Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI*. Penguin.

Harari, Yuval Noah. 2024. *Nexus: A Brief History of Information Networks from the Stone Age to AI*. Random House.

Hardman, Philippa. 2025. "The AI Illusion in L&D." *Dr Phil's Newsletter, Powered by DOMSTM AI* (blog), January 31. <https://drphilppahardman.substack.com/p/the-ai-illusion-in-l-and-d>.

Haskins, Caroline, and Lauren Goode. 2025. "Grok Is Spewing Antisemitic Garbage on X." *Wired*, September 15. <https://www.wired.com/story/grok-antisemitic-posts-x-xai/>.

Hatch, S. Gabe, Zachary T. Goodman, Laura Vowels, et al. 2025. "When ELIZA Meets Therapists: A Turing Test for the Heart and Mind."

*PLOS Mental Health* 2 (2): e0000145. <https://doi.org/10.1371/journal.pmen.0000145>.

Hattie, John. 2010. *Visible Learning*. Routledge.

Hector, Hamish. 2025. "Are You Polite to ChatGPT? Here's Where You Rank among AI Chatbot Users." TechRadar, February 20. <https://www.techradar.com/computing/artificial-intelligence/are-you-polite-to-chatgpt-heres-where-you-rank-among-ai-chatbot-users>.

Heinz, Michael V., Daniel M. Mackin, Brianna M. Trudeau, et al. 2025. "Randomized Trial of a Generative AI Chatbot for Mental Health Treatment." *New England Journal of Medicine AI* 2 (4): A1oa2400802. <https://doi.org/10.1056/A1oa2400802>.

Hendrycks, Dan, Collin Burns, Steven Basart, et al. 2021. "Measuring Massive Multitask Language Understanding." Preprint, arXiv:2009.0330, January 12, 2021. <https://doi.org/10.48550/arXiv.2009.03300>.

Henrich, Joseph, Steven J. Heine, and Ara Norenzayan. 2010. "The Weirdest People in the World?" *Behavioral and Brain Sciences* 33 (2–3): 61–83. <https://doi.org/10.1017/S0140525X0999152X>.

Heritage, Stuart. 2025. "'I Felt Pure, Unconditional Love': The People Who Marry Their AI Chatbots." *The Guardian*, July 12. <https://www.theguardian.com/tv-and-radio/2025/jul/12/i-felt-pure-unconditional-love-the-people-who-marry-their-ai-chatbots>.

Hill, Heather C., John P. Papay, and Nathaniel Schwartz. 2022. *Dispelling the Myths: What the Research Says about Teacher Professional Learning*. February 15, 10. <https://annenbergbrown.edu/sites/default/files/rppl-dispelling-myths.pdf>.

Hill, Kashmir. 2025. "OpenAI Plans to Add Safeguards to ChatGPT for Teens and Others in Distress." *The New York Times*, September 2. <https://www.nytimes.com/2025/09/02/technology/>

[personaltech/chatgpt-parental-controls-openai.html](https://personaltech.com/chatgpt-parental-controls-openai.html).

Ho, Annabell, Jeff Hancock, and Adam S. Miner. 2018. "Psychological, Relational, and Emotional Effects of Self-Disclosure After Conversations with a Chatbot." *The Journal of Communication* 68 (4): 712–33. <https://doi.org/10.1093/joc/jqy026>.

Holzer, Andreas, and Martin Daumiller. 2025. "Building Trust in the Classroom: Perspectives from Students and Teachers." *European Journal of Psychology of Education* 40 (62). <https://doi.org/10.1007/s10212-025-00961-7>.

Hord, Shirley M., William L. Rutherford, Leslie Huling, and Gene E. Hall. 2006. *Taking Charge of Change*. Southwest Educational Development Laboratory.

Hou, Irene, Sophia Metille, Zhuo Li, Owen Man, Cynthia Zastudil, and Stephen MacNeil. 2024. "The Effects of Generative AI on Computing Students' Help-Seeking Preferences." *Proceedings of the 26th Australasian Computing Education Conference (ACE '24)*. January 29. Association for Computing Machinery, New York, NY, USA, 39–48. <https://doi.org/10.1145/3636243.3636248>.

Howdle, Dan. 2023. "The Cost of 1GB Of Mobile Data in 237 Countries." Bestbroadbanddeals. <https://bestbroadbanddeals.co.uk/mobiles/world-wide-data-pricing/>.

HP. 2025. "AI and Education: A Helpful Guide for." HP, May 7. <https://www.hp.com/nz-en/shop/tech-takes/post/ai-education-guide-for-parents>.

Hradecky, Alex, and Abdulhakeem Abdulkareem. 2025. "African Entrepreneurs Are Redefining How AI Can Drive Sustainable Development." UNDP Digital, AI and Innovation Hub, September 2. <https://www.undp.org/digital/blog/african-entrepreneurs-are-redefining-how-ai-can-drive-sustainable-development>.

Huckins, Grace. 2025. "OpenAI Has Finally Released Open-Weight Language Models." *MIT Technology Review*, October 14. <https://www.technologyreview.com/2025/08/05/1121092/openai-has-finally-released-open-weight-language-models/>.

Huxley, Aldous. 1932. *Brave New World*. Harper Perennial.

Ihichr, Adel, Omar Oustous, Younes El Bouzekri El Idrissi, and Ayoub Ait Lahcen. 2024. "A Systematic Review on Assessment in Adaptive Learning: Theories, Algorithms and Techniques." *International Journal of Advanced Computer Science and Applications* 15 (7). <https://doi.org/10.14569/IJACSA.2024.0150785>.

Imagine Learning. 2025. "Imagine Learning." Imagine Learning. <https://www.imaginelearning.com/pdf-viewer/>.

Immordino-Yang, Mary Helen. 2022. "Purpose, Feelings, Deep Thinking, and Relationships Drive Brain Development." Kappan Online, March 3. <https://kappanonline.org/purpose-feelings-deep-thinking-and-relationships-brain-development-immordino-yang/>.

Immordino-Yang, Mary Helen. 2025. "How Transcendent Thinking Boosts Teen Brains and Enhances Life." *Scientific American*, January 21. <https://www.scientificamerican.com/article/transcendent-thinking-boosts-teen-brains-in-ways-that-enhance-life/>.

Immordino-Yang, Mary Helen, and Antonio Damasio. 2007. "We Feel, Therefore We Learn: The Relevance of Affective and Social Neuroscience to Education." *Mind, Brain, and Education* 1 (1): 3–10. <https://doi.org/10.1111/j.1751-228X.2007.00004.x>.

Independent Communications Authority of South Africa. 2021. *Electronic Communications Act: Licensing Process for International Mobile Applications in Respect of the Provision of Mobile Broadband Wireless Access Services: Invitation to*

- Apply. [https://www.gov.za/sites/default/files/gcis\\_document/202112/45628gen717.pdf](https://www.gov.za/sites/default/files/gcis_document/202112/45628gen717.pdf).
- Infocomm Media Development Authority. 2025. "A Guide to Cyber Safety for Kids at Every Age." Infocomm Media Development Authority. <https://www.imda.gov.sg/resources/blog/blog-articles/2025/09/cyber-safety-for-kids-at-every-age>.
- Ingersoll, Richard, Lennon Audrain, and Mary Laski. n.d. "Team-Based Staffing, Teacher Authority, and Teacher Turnover." <https://crpe.org/team-based-staffing-teacher-authority-and-teacher-turnover/>.
- Ingersoll, Richard M., and Richard Audrain. 2025. "The Teaching Model That's Keeping Educators in Schools." *ASCD* 82 (7). <https://www.ascd.org/el/articles/the-teaching-model-thats-keeping-educators-in-schools>.
- International Society for Technology in Education (ISTE). 2024. "Evolving Teacher Education in an AI World." ISTE. <https://info.iste.org/evolvingteacher-ed-ai>.
- International Telecommunication Union (ITU). 2023a. "Economy Classifications." ITU. <https://www.itu.int:443/en/ITU-D/Statistics/pages/definitions/regions.aspx>.
- International Telecommunication Union. 2023b. "Measuring Digital Development: Facts and Figures." ITU. [https://www.itu.int/hub/publication/d-ind-ict\\_mdd-2023-1/](https://www.itu.int/hub/publication/d-ind-ict_mdd-2023-1/).
- International Telecommunication Union. 2024. "Facts and Figures 2024—Mobile Phone Ownership." ITU. <https://www.itu.int/itu-d/reports/statistics/2024/11/10/ff24-mobile-phone-ownership>.
- Internet Matters Team. 2025. "Me, Myself and AI Chatbot Research." Internet Matters. <https://www.internetmatters.org/hub/research/me-myself-and-ai-chatbot-research/>.
- Ipsos. 2024. "Common Sense Media Ipsos Generative AI Survey of Parents and Teens." Ipsos. <https://www.ipsos.com/sites/default/files/ct/news/documents/2024-09/Common%20Sense%20Media%20Ipsos%20Generative%20AI%20Survey%20of%20Parents%20and%20Teens%2009192024.pdf>.
- ISTE (International Society for Technology in Education). 2025. "Artificial Intelligence in Education." ISTE. <https://iste.org/ai>.
- iThrive Games. 2025. "Home." iThrive Games, October 26. <https://ithrivegames.org>.
- Iyengar, Sunil. 2024. "Federal Data on Reading for Pleasure: All Signs Show a Slump." *National Endowment for the Arts* (blog), October 3. <https://www.arts.gov/stories/blog/2024/federal-data-reading-pleasure-all-signs-show-slump>.
- JA Worldwide. n.d. "JA Boost: Where AI Meets the Future of Learning." JA Worldwide, accessed December 15, 2025. <https://www.jaworldwide.org/jaboost>.
- Jain, Ronak, and Samuel Stemper. 2024. "3G Internet and Human Capital Development." University of Zurich Department of Economics Working Paper no. 453. <https://www.ictworks.org/wp-content/uploads/2025/07/3G-use-reduces-test-score.pdf>.
- Jesuit, David K., and Brian Endless. 2018. "Model United Nations and Experiential Learning: An Assessment of Changes in Knowledge and Attitudes." *Journal of Social Studies Education Research* 9 (4): 198–213.
- Jiang, Yuan-Hao, Yuang Wei, Xiaobao Shao, Rui Jia, Yizhou Zhou, and Zi-Wei Chen. 2025. "Generative AI in Personalized Learning: Development Trajectory, Educational Applications, and Future Education." Paper presented at the 36th annual conference of the Society for Information Technology and Teacher Education (SITE 2025), Orlando, Florida,

March 17, 2025. [https://www.researchgate.net/publication/389826421\\_Generative\\_AI\\_in\\_Personalized\\_Learning\\_Development\\_Trajectory\\_Educational\\_Applications\\_and\\_Future\\_Education](https://www.researchgate.net/publication/389826421_Generative_AI_in_Personalized_Learning_Development_Trajectory_Educational_Applications_and_Future_Education).

Johansen, Vegard. 2018. *Innovation Cluster for Entrepreneurship Education*. Østlandsforskning/ Eastern Norway Research Institute. [https://www.researchgate.net/publication/327623618\\_Innovation\\_Cluster\\_for\\_Entrepreneurship\\_Education#fullTextFileContent](https://www.researchgate.net/publication/327623618_Innovation_Cluster_for_Entrepreneurship_Education#fullTextFileContent).

Johinke, Rebecca, Robert Cummings, and Frances Di Lauro. 2023. "Reclaiming the Technology of Higher Education for Teaching Digital Writing in a Post-Pandemic World." *Journal of University Teaching & Learning Practice* 20 (2). <https://doi.org/10.53761/1.20.02.01>.

Johnson, Arianna. 2023. "ChatGPT in Schools: Here's Where It's Banned—And How It Could Potentially Help Students." *Forbes*, January 18. <https://www.forbes.com/sites/ariannajohnson/2023/01/18/chatgpt-in-schools-heres-where-its-banned-and-how-it-could-potentially-help-students/>.

Johnson, Jaivarsini. 2024. "Effect of Emotions on Learning, Memory, and Disorders Associated with the Changes in Expression Levels: A Narrative Review." *Brain Circulation* 10 (2): 134–44. [https://doi.org/10.4103/bc.bc\\_86\\_23](https://doi.org/10.4103/bc.bc_86_23).

Johnson, Matt. 2021. "The Rising Loneliness Economy." *Psychology Today*, December 16. <https://www.psychologytoday.com/us/blog/mind-brain-and-value/202112/the-rising-loneliness-economy>.

Joon-hyun, Moon. 2025. "AI Chats Feel 'Emotionally Meaningful,' Say about 40% of Young South Koreans in Survey." *The Korea Herald*, February 26. <https://www.koreaherald.com/article/10429545>.

Kai, Ng Wei. 2022. "NIE to Train Teachers in Using AI in Classroom, Invest in Research." *The Straits Times*, May 25. <https://www.straitstimes.com/>

[singapore/parenting-education/nie-to-train-teachers-in-using-ai-in-classroom-invest-in-research](https://www.straitstimes.com/singapore/parenting-education/nie-to-train-teachers-in-using-ai-in-classroom-invest-in-research).

Kaliisa, Rogers, Kamila Misiejuk, Sonsoles López-Pernas, and Mohammed Saqr. 2025. "How Does Artificial Intelligence Compare to Human Feedback? A Meta-Analysis of Performance, Feedback Perception, and Learning Dispositions." *Educational Psychology*, September 24, 1–32. <https://doi.org/10.1080/01443410.2025.2553639>.

Kang, Cecilia. 2023. "Sam Altman, ChatGPT Creator and OpenAI CEO, Urges Senate for AI Regulation." *The New York Times*, May 16. <https://www.nytimes.com/2023/05/16/technology/openai-altman-artificial-intelligence-regulation.html>.

Kang, Cecilia. 2025. "A.I.-Generated Images of Child Sexual Abuse Are Flooding the Internet." *The New York Times*, July 10. <https://www.nytimes.com/2025/07/10/technology/ai-csam-child-sexual-abuse.html>.

Kannan, Prabha. 2024. "How Harmful Are AI's Biases on Diverse Student Populations?" Stanford University Human-Centered Artificial Intelligence, October 3. <https://hai.stanford.edu/news/how-harmful-are-ais-biases-on-diverse-student-populations>.

Kasneci, Enkelejda, Kathrin Sessler, Stefan Küchemann, et al. 2023. "ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education." *Learning and Individual Differences* 103 (April): 102274. <https://doi.org/10.1016/j.lindif.2023.102274>.

Kaufman, Julia H., Ashley Woo, Joshua Eagan, Sabrina Lee, and Emma B. Kassan. 2025. *Uneven Adoption of Artificial Intelligence Tools Among U.S. Teachers and Principals in the 2023–2024 School Year*. RAND, February 11. [https://www.rand.org/pubs/research\\_reports/RRA134-25.html](https://www.rand.org/pubs/research_reports/RRA134-25.html).

Kelly, Rhea. 2025. "New National Academy for AI Instruction to Provide Free AI Training for Educators." *Technological Horizons in Education*

(THE) Journal, July 9. <https://thejournal.com/Articles/2025/07/09/New-National-Academy-for-AI-Instruction-to-Provide-Free-AI-Training-for-Educators.aspx>.

KERIS (Korea Education and Research Information Service). 2025. "Digital Education for All: KERIS." KERIS, September 29. <https://www.keris.or.kr/eng/main.do>.

Kestin, Greg, Kelly Miller, Anna Klaes, Timothy Milbourne, and Gregorio Ponti. 2025. "AI Tutoring Outperforms In-Class Active Learning: An RCT Introducing a Novel Research-Based Design in an Authentic Educational Setting." *Scientific Reports* 15 (1): 17458. <https://doi.org/10.1038/s41598-025-97652-6>.

Khan Academy. 2025. "Khanmigo for Learners: Always-Available Tutor, Powered by AI." Khan Academy, September 13. <https://www.khanmigo.ai/learners>.

Khodyakov, Dmitry. 2023. "Generating Evidence Using the Delphi Method." RAND, October 17. <https://www.rand.org/pubs/commentary/2023/10/generating-evidence-using-the-delphi-method.html>.

Kleiman, Glenn M., and H. Alix Gallagher. 2024. "State Education Policy and the New Artificial Intelligence." National Association of State Boards of Education (NASBE). <https://www.nasbe.org/state-education-policy-and-the-new-artificial-intelligence/>.

Klein, Gary. 2007. "Performing a Project Premortem." *Harvard Business Review* 85 (9): 18–19.

Klein, Gary. 2021. "The Pre-Mortem Method." *Psychology Today*, January 14. <https://www.psychologytoday.com/us/blog/seeing-what-others-dont/202101/the-pre-mortem-method>.

Klein, Lauren, Meredith Martin, André Brock, et al. 2025. "Provocations from the Humanities

for Generative AI Research." Preprint, arXiv:2502.19190, February 26, 2025. <https://doi.org/10.48550/arXiv.2502.19190>.

Kosmyna, Nataliya, Eugene Hauptmann, Ye Tong Yuan, et al. 2025. "Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task." Preprint, arXiv:2506.08872, June 10, 2025. <https://doi.org/10.48550/arXiv.2506.08872>.

Kranzberg, Melvin. 1986. "Technology and History: 'Kranzberg's Laws.'" *Technology and Culture* 27 (3): 544–60. <https://doi.org/10.2307/3105385>.

Krishnamurti, Jiddu. 1953. *Education and the Significance of Life*. Harper.

Kuerban, Yushan, Solomon Sunday Oyelere, and Ismaila Temitayo Sanusi. 2025. "ReadSmart: Generative AI and Augmented Reality Solution for Supporting Students with Dyslexia Learning Disabilities." *International Journal of Technology in Education and Science* 9 (1): 159–76. <https://doi.org/10.46328/ijtes.599>.

Kulik, James A., and J. D. Fletcher. 2016. "Effectiveness of Intelligent Tutoring Systems: A Meta-Analytic Review." *Review of Educational Research* 86 (1): 42–78. <https://doi.org/10.3102/0034654315581420>.

Kumar, Ajay, Thomas H. Davenport, and Randy Bean. 2025. "The Case for Using Small Language Models." *Harvard Business Review*, October 14. <https://hbr.org/2025/09/the-case-for-using-small-language-models>.

Kumayama, Ken D., Stuart D. Levi, William E. Ridgway, et al. 2025. "Landmark California AI Safety Legislation May Serve as a Model for Other States in the Absence of Federal Standards." Skadden, October 2. <https://www.skadden.com/insights/publications/2025/10/landmark-california-ai-safety-legislation>.

- Kurian, Nomisha. 2024. "'No, Alexa, No!': Designing Child-Safe AI and Protecting Children from the Risks of the 'Empathy Gap' in Large Language Models." *Learning, Media and Technology*, July 10. <https://doi.org/10.1080/17439884.2024.2367052>.
- Laird, Elizabeth, Maddy Dwyer, and Hannah Quay-de la Vallee. 2025. "Hand in Hand: Schools' Embrace of AI Connected to Increased Risks to Students." Center for Democracy and Technology, October 8. <https://cdt.org/insights/hand-in-hand-schools-embrace-of-ai-connected-to-increased-risks-to-students/>.
- Lake, Robin. 2023. "Shockwaves & Innovations: How Nations Worldwide Are Dealing with AI in Education." *The 74 Million*, August 7. <https://www.the74million.org/article/shockwaves-innovations-how-nations-world-wide-are-dealing-with-ai-in-education/>.
- Lanier, Jaron. 2025. "Your A.I. Lover Will Change You." *The New Yorker*, March 22. <https://www.newyorker.com/culture/the-weekend-essay/your-ai-lover-will-change-you>.
- Lanz, Jose Antonio. 2025. "AI's Premium Price Tags Are Creating a New Digital Divide." *Decrypt*, February 23. <https://decrypt.co/306117/ai-premium-price-tags-new-digital-divide>.
- Latham and Watkins. 2025. "California Assumes Role as Lead US Regulator of AI." Latham and Watkins LLP, October 15. <https://www.lw.com/en/insights/california-assumes-role-as-lead-us-regulator-of-ai>.
- Learn and Work. 2025. "AI Literacy Framework for Primary & Secondary Education—OECD-EC." Learn & Work Ecosystem Library. <https://learnworkecosystemlibrary.com/initiatives/ai-literacy-framework-for-primary-secondary-education-oecd-ec/>.
- Learning Accelerator. 2025. "Driving EdTech Systems: Student Data Privacy." The Learning Accelerator, December 1. <https://practices.learningaccelerator.org/strategies/driving-ed-tech-systems-student-data-privacy>.
- Learning Equality. 2025. "Home Page." Learning Equality, October 15. <http://learningequality.org/>.
- Learning Policy Institute. 2024. "2024 Update: What's the Cost of Teacher Turnover?" Learning Policy Institute, September 17. <https://learningpolicyinstitute.org/product/2024-whats-cost-teacher-turnover>.
- LearnLab. 2025. "Introduction to Learning Engineering." LearnLab, October 26. <https://learnlab.org/introduction-to-learning-engineering/>.
- LearnLM Team, Google, Alicia Martín, et al. 2025. "Towards an AI-Augmented Textbook." Preprint, arXiv:2509.13348, September 30, 2025. <https://doi.org/10.48550/arXiv.2509.13348>.
- Lee, Hao-Ping, Sarkar Advait, Lev Tankelevitch, et al. 2025. "The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects from a Survey of Knowledge Workers." *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* Article 1121 (April): 1–22. <https://doi.org/10.1145/3706598.3713778>.
- Lelièvre, Maxime, Amy Waldock, Meng Liu, et al. 2025. "Benchmarking the Pedagogical Knowledge of Large Language Models." Preprint, arXiv:2506.18710. Version 1. June 23, 2025. <https://doi.org/10.48550/arXiv.2506.18710>.
- Lembke, Anna. 2023. "Applying the Bradford Hill Criteria to Social Media Use and Adolescent Mental Health." *After Babel* (blog), March 9. <https://www.afterbabel.com/p/bradford-hill-social-media>.
- Lesan. 2025. "Lesan AI." Lesan, October 26. <https://lesan.ai/about.html>.
- Letort, Brian, and Kadri Linsak-Goode. 2025. "What Is Sovereign AI?." Digital Realty, October 10.

<https://www.digitalrealty.com/resources/articles/what-is-sovereign-ai?t=1760120896645?latest>.

Levia, Benjamin, Ana T. Ribeiro, Chris Agnew, and Susanna Loeb. 2025. *How K-12 Educators Are Actually Engaging with AI: Early Findings from an EdTech Platform*. Stanford Scale. <https://scale.stanford.edu/research-in-action/how-k12-educators-engaging-with-ai>.

Levy, Steven. 2025. "Trump's Anti-Bias AI Order Is Just More Bias." *WIRED*, July 25. <https://www.wired.com/story/trump-ai-order-bias-openai-google/>.

Lin, Zhiqiang, Huan Sun, and Ness Shroff. 2025. "AI Safety vs. AI Security: Demystifying the Distinction and Boundaries." Preprint, arXiv:2506.18932, Version 1. June 21, 2025. <https://doi.org/10.48550/arXiv.2506.18932>.

Long, Duri, and Brian Magerko. 2020. "What Is AI Literacy? Competencies and Design Considerations." *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, April 23, 1–16. <https://doi.org/10.1145/3313831.3376727>.

Lunabase.ai. 2025. "The Evolution of AI Language Models: From ChatGPT to GPT-5 and Beyond." Luna, October 21. <https://lunabase.ai/blog/the-evolution-of-ai-language-models-from-chatgpt-to-gpt-5-and-beyond>.

Luo, Heng. 2023. "Editorial: Advances in Multimodal Learning: Pedagogies, Technologies, and Analytics." *Frontiers in Psychology* 14 (October): 1286092. <https://doi.org/10.3389/fpsyg.2023.1286092>.

Luo, Heng, Gege Li, Qinna Feng, Yuqin Yang, and Mingzhang Zuo. 2021. "Virtual Reality in K-12 and Higher Education: A Systematic Review of the Literature from 2000 to 2019." *Journal of Computer Assisted Learning* 37 (March): 887–901. <https://doi.org/10.1111/jcal.12538>.

Lyons-Cunha, Jenny. 2025. "What Are AI Companions?" *Built In*, November 19. <https://builtin.com/artificial-intelligence/ai-companions>.

Macleod, Ewen. 2021. "Data and Evidence on Forced Displacement: Reflections on Progress and Challenges." *Force Migration Review* 66 (Mental Health and Psychosocial Support): 46–50.

Madden, Mary, Angela Calvin, Alexa Hasse, and Amanda Lenhart. 2024. *The Dawn of the AI Era: Teens, Parents, and the Adoption of Generative AI at Home and School*. Common Sense. [https://www.common SenseMedia.org/sites/default/files/research/report/2024-the-dawn-of-the-ai-era\\_final-release-for-web.pdf](https://www.common SenseMedia.org/sites/default/files/research/report/2024-the-dawn-of-the-ai-era_final-release-for-web.pdf).

Mah, Christopher, Tan, Mei, Phalen, Lena, Sparks, Alexa, and Demszky, Dorottya. 2025. "From Sentence-Corrections to Deeper Dialogue: Qualitative Insights from LLM and Teacher Feedback on Student Writing." *EdWorkingPaper* No. 25–1193 (May). <https://doi.org/10.26300/P397-2P46>.

Mahari, Robert, and Pat Pataranutaporn. 2025. "Addictive Intelligence: Understanding Psychological, Legal, and Technical Dimensions of AI Companionship." *MIT Case Studies in Social and Ethical Responsibilities of Computing* (Winter 2025). <https://doi.org/10.21428/2c646de5.2877155b>.

Mahari, Robert, and Alex Pentland. 2024. "Regulation by Design: A New Paradigm for Regulating AI Systems." *Digital Single Market and Artificial Intelligence: AI Act and Intellectual Property in the Digital Transition* (17): 431–41. <https://doi.org/10.53136/979122181150627>.

Major, Louis, and Gill A. Francis. 2020. *Technology-Supported Personalised Learning: A Rapid Evidence Review*. EdTech Hub/Zenodo. <https://doi.org/10.5281/zenodo.4556925>.

Major, Louis, Gill A. Francis, and Maria Tsapali. 2021. "The Effectiveness of Technology-Supported Personalised Learning in Low- and Middle-Income

Countries: A Meta-Analysis." *British Journal of Educational Technology* 52: 1935–64. <https://doi.org/10.1111/bjet.13116>.

Malek Ash, Andrea. 2025. "Three in 10 Teachers Use AI Weekly, Saving Six Weeks a Year." Gallup.com, June 24. <https://news.gallup.com/poll/691967/three-teachers-weekly-saving-six-weeks-year.aspx>.

Maples, Bethanie, Merve Cerit, Aditya Vishwanath, and Roy Pea. 2024. "Loneliness and Suicide Mitigation for Students Using GPT3-Enabled Chatbots." *NPJ Mental Health Research* 3 (1): 4. <https://doi.org/10.1038/s44184-023-00047-6>.

Maples, Bethanie, Roy D. Upadyaya, and David Markowitz. 2023. "Learning from Intelligent Social Agents as Social and Intellectual Mirrors." In *AI in Learning: Designing the Future*, edited by Hannele Niemi, Roy D. Pea, and Yu Lu. Springer International Publishing. <https://doi.org/10.1007/978-3-031-09687-7>.

Market.us. 2025. *Global AI Companion App Market Size, Share, Statistics Analysis Report*. Market.us No. 136977. <https://market.us/report/ai-companion-app-market/>.

Marton, F., and R. Säljö. 1976. "On Qualitative Differences in Learning: I—Outcome and Process." *British Journal of Educational Psychology* 46 (1): 4–11. <https://doi.org/10.1111/j.2044-8279.1976.tb02980.x>.

Marzuki, Utami Widiati, Diyenti Rusdin, Darwin, and Inda Indrawati. 2023. "The Impact of AI Writing Tools on the Content and Organization of Students' Writing: EFL Teachers' Perspective." *Cogent Education* 10 (2): 2236469. <https://doi.org/10.1080/2331186X.2023.2236469>.

Masakhane. 2025. "Masakhane: A Grassroots NLP Community for Africa, by Africans." Masakhane, October 26. <https://www.masakhane.io/>.

Masters, Geoff. 2021. "Computer Adaptive Testing—Challenging Traditional Thinking." *Teacher Magazine*, June 21. [https://www.teachermagazine.com/au\\_en/articles/computer-adaptive-testing-challenging-traditional-thinking](https://www.teachermagazine.com/au_en/articles/computer-adaptive-testing-challenging-traditional-thinking).

McCabe, David. 2025. "Regulators Are Digging into A.I. Chatbots and Child Safety." *The New York Times*, September 11. <https://www.nytimes.com/2025/09/11/technology/google-meta-chatgpt-ai-chatbots.html>.

McCrae, Phil. 2025. "Artificial Intelligence in Education: Anticipating the Disruptions and Mapping the Future." Alberta Teachers Association Social Studies Specialist Council Provincial Event, May 8.

McKinsey & Company. 2024. "What Are AI Guardrails?" McKinsey, November 14. <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-are-ai-guardrails>.

McLennan, April. 2025. "AI Chatbots Accused of Encouraging Teen Suicide as Experts Sound Alarm." *RNZ*, August 15. <https://www.rnz.co.nz/news/world/570138/ai-chatbots-accused-of-encouraging-teen-suicide-as-experts-sound-alarm>.

Mehra, Samiksha. 2020. "How India Is Integrating AI in the New Education Policy." *IndiaAI*, August 3. <https://indiaai.gov.in/article/how-india-is-integrating-ai-in-the-new-education-policy>.

Mehta, Arun C. 2020. "Online Education, Elementary Education." *Education for All in India*, October 18. <https://educationforallinindia.com/online-education/>.

Mehta, Ivan. 2025. "OpenAI Says ChatGPT Is on Track to Reach 700M Weekly Users." *TechCrunch*, August 4. <https://techcrunch.com/2025/08/04/openai-says-chatgpt-is-on-track-to-reach-700m-weekly-users/>.

Meniado, Joel C., Duong Thi Thu Huyen, Nopparat Panyadilokpong, and Pannaphatt Lertkomolwit.

2024. "Using ChatGPT for Second Language Writing: Experiences and Perceptions of EFL Learners in Thailand and Vietnam." *Computers and Education: Artificial Intelligence* 7 (December): 100313. <https://doi.org/10.1016/j.caeai.2024.100313>.
- Merod, Anna. 2025. "Ransomware Attacks Surge 69% across Global Education Sector." K-12 Dive, April 14. <https://www.k12dive.com/news/ransomware-attacks-surge-69-across-global-education-sector/745153/>.
- Meyer, Jennifer, Thorben Jansen, Ronja Schiller, et al. 2024. "Using LLMs to Bring Evidence-Based Feedback into the Classroom: AI-Generated Feedback Increases Secondary Students' Text Revision, Motivation, and Positive Emotions." *Computers and Education: Artificial Intelligence* 6 (June): 100199. <https://doi.org/10.1016/j.caeai.2023.100199>.
- Miao, Fengchun, Wayne Holmes, Ronghuai Huang, and Hui Zhang. 2021. AI and Education: Guidance for Policy-Makers. UNESCO. <https://doi.org/10.54675/PCSP7350>.
- Miao, Fengchun, and Wayne Holmes. 2023. *Guidance for Generative AI in Education and Research*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000386693>.
- Microsoft. 2025a. *2025 AI in Education: A Microsoft Special Report*. Microsoft. <https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/bade/documents/products-and-services/en-us/education/2025-Microsoft-AI-in-Education-Report.pdf>.
- Microsoft. 2025b. "What Is a Chatbot?" Microsoft. <https://www.microsoft.com/en/microsoft-copilot/copilot-101/what-is-a-chatbot>.
- Microsoft Source. 2025a. "AFT to Launch National Academy for AI Instruction with Microsoft, OpenAI, Anthropic and United Federation of Teachers." Microsoft Source, July 8. [al-academy-for-ai-instruction-with-microsoft-openai-anthropic-and-united-federation-of-teachers/](https://news.microsoft.com/source/2025/07/08/aft-to-launch-nation-al-academy-for-ai-instruction-with-microsoft-openai-anthropic-and-united-federation-of-teachers/).
- Microsoft Source. 2025b. "Pearson and Microsoft Announce Multiyear Partnership to Transform the Future of Learning and Work with AI." Microsoft Source, January 14. <https://news.microsoft.com/source/2025/01/14/pearson-and-microsoft-announce-multiyear-partnership-to-transform-the-future-of-learning-and-work-with-ai/>.
- Milmo, Dan. 2022. "Molly Russell: How Family Are Helping Shift Narrative on Online Safety." *The Guardian*, September 30. <https://www.theguardian.com/technology/2022/sep/30/molly-russell-family-helping-shift-online-safety-regulation>.
- Mishra, Punya. 2024. "AI'nt Fair: Why AI May Make Learning Gaps Wider." *Punya Mishra* (blog), November 19. <https://punyamishra.com/2024/11/19/aint-fair-why-ai-may-make-learning-gaps-wider/>.
- Mishra, Punya. 2025. "Engineered for Attachment: The Hidden Psychology of AI Companions." *Punya Mishra* (blog), August 1. <https://punyamishra.com/2025/05/02/engineered-for-attachment-the-hidden-psychology-of-ai-companions/>.
- Mitchell, Deborah J., J. Edward Russo, and Nancy Pennington. 1989. "Back to the Future: Temporal Perspective in the Explanation of Events." *Journal of Behavioral Decision Making* 2 (1): 25. <https://doi.org/10.1002/BDM.3960020103>.
- Mollick, Ethan, Lilach Mollick, Natalie Bach, L. J. Ciccarelli, Ben Przystanski, and Daniel Ravipinto. 2024. "AI Agents and Education: Simulated Practice at Scale." Preprint, arXiv:2407.12796, June 20, 2024. <https://doi.org/10.48550/arXiv.2407.12796>.
- Moon, Kibum, Adam E. Green, and Kostadin Kushlev. 2025a. "Homogenizing Effect of Large Language Models (LLMs) on Creative Diversity: An Empirical Comparison of Human and ChatGPT Writing." *Computers in Human Behavior: Artificial*

Humans 6 (December): 100207. <https://doi.org/10.1016/j.chbah.2025.100207>.

Moon, Kibum, Kostadin Kushlev, Andrew Bank, Indre Viskontas, and Adam Green. 2025b. "LLM-Era College Admissions Essays Exhibit Paradoxical Semantic Trends." Preprint, PsyArXiv, Jsz58\_v4. October 22, 2025. [https://doi.org/10.31234/osf.io/jsz58\\_v4](https://doi.org/10.31234/osf.io/jsz58_v4).

Moore, Steven, Lydia Eckstein, Christine Kwon, and John Stamper. 2025. "Generative AI in Instructional Design Education: Effects on Novice Microlesson Quality." In *Artificial Intelligence in Education*, edited by Alexandra I. Cristea, Erin Walker, Yu Lu, Olga C. Santos, and Seiji Isotani. Springer Nature Switzerland. [https://doi.org/10.1007/978-3-031-98459-4\\_35](https://doi.org/10.1007/978-3-031-98459-4_35).

Moorosi, Nyalleng. 2024. "Better Data Sets Won't Solve the Problem—We Need AI for Africa to Be Developed in Africa." *Nature* 636 (8042): 276–276. <https://doi.org/10.1038/d41586-024-03988-w>.

Mormando, Samuel. 2025. "A Stoplight Model for Guiding Student AI Usage." *Edutopia*, December 1. <https://www.edutopia.org/article/creating-ai-usage-guidelines-students/>.

Mosseri, Adam, and Alexandr Wang. 2025. "Our Approach to Teen AI Safety: Empowering Parents, Protecting Teens." *Meta*, October 17. <https://about.fb.com/news/2025/10/teen-ai-safety-approach/>.

Mounk, Yasha. 2024. "Freya India on How to Free the Anxious Generation." *Persuasion* (blog), November 15. <https://www.persuasion.community/p/freya-india-on-how-to-free-the-anxious>.

Mozilla Foundation. 2025. "About Our Methodology." Mozilla Foundation, September 10. <https://www.mozillafoundation.org/en/privacynotincluded/>.

Muncey, Nikki. 2025. "How Teachers Can Use AI Chatbots in History Class." *SchoolAI*, August 15.

<https://schoolai.com/blog/engaging-in-history-class-using-ai-chatbots-as-a-teaching-tool>.

Myers, Steven Lee, and Stuart A. Thompson. 2025. "Right-Wing Chatbots Turbocharge America's Political and Cultural Wars." *The New York Times*, November 4. <https://www.nytimes.com/2025/11/04/business/right-wing-chatbots-gab-arya-chatgpt-gemini.html>.

Nakashima, Ryan. 2018. "Google Tracks Your Movements, Like It or Not." *AP News*, August 13. <https://apnews.com/article/828aefab64d4411ba-c257a07c1af0ecb>.

Natale, Simone. 2021. *Deceitful Media: Artificial Intelligence and Social Life after the Turing Test*. Oxford University Press. <https://doi.org/10.1093/oso/9780190080365.001.0001>.

Natale, Simone, and Iliana Depounti. 2024. "Artificial Sociality." *Human-Machine Communication* 7: 83–98. <https://doi.org/10.30658/hmc.7.5>.

National Academies of Sciences, Engineering, and Medicine. 2018. *How People Learn II: The Science and Practice of Learning*. National Academies Press. <https://doi.org/10.17226/24783>.

National Center for Education Statistics. 2024. "NAEP Long-Term Trend Assessment Results: Reading and Mathematics." *The Nation's Report Card*. <https://www.nationsreportcard.gov/highlights/ltt/2023/>.

National Academies of Sciences, Engineering, and Medicine. 2018. *How People Learn II: The Science and Practice of Learning*. National Academies Press. <https://doi.org/10.17226/24783>.

National Commission on Teaching and America's Future. 2003. *The High Cost of Teacher Turnover*. <https://eric.ed.gov/?id=ED498001>.

National Education Association (NEA). 2025. "Vetting AI Resources." NEA, June 20. <https://www.>

[nea.org/professional-excellence/student-engagement/tools-tips/vetting-ai-resources](https://www.nea.org/professional-excellence/student-engagement/tools-tips/vetting-ai-resources).

National Scientific Council on the Developing Child. 2004. "Young Children Develop in an Environment of Relationships" Working Paper No. 1. <https://developingchild.harvard.edu/resources/working-paper/wp1/>.

Newman, Lily Hay. 2021. "iOS 14.5 Lets You Stop Ads from Tracking You—So Do It." *WIRED*, April 26. <https://www.wired.com/story/ios-app-tracking-transparency-advertising/>.

Newman, Lily Hay. 2025. "Vibe Coding Is the New Open Source—in the Worst Way Possible." *WIRED*, October 6. <https://www.wired.com/story/vibe-coding-is-the-new-open-source/>.

Ng, Davy Tsz Kit, Chee Wei Tan, and Jac Ka Lok Leung. 2024. "Empowering Student Self-Regulated Learning and Science Education through ChatGPT: A Pioneering Pilot Study." *British Journal of Educational Technology* 55 (4): 1328–53. <https://doi.org/10.1111/bjjet.13454>.

Nickow, Andre, Philip Oreopoulos, and Vincent Quan. 2020. "The Impressive Effects of Tutoring on PreK–12 Learning: A Systematic Review and Meta-Analysis of the Experimental Evidence." National Bureau of Economic Research Working Paper No. 27476, July. <https://doi.org/10.3386/w27476>.

Niemi, Hannele, Roy D. Pea, and Yu Lu, eds. 2023. *AI in Learning: Designing the Future*. Springer Nature. <https://doi.org/10.1007/978-3-031-09687-7>.

Niloy, Ahnaf Chowdhury, Salma Akter, Nayeema Sultana, Jakia Sultana, and Sayed Imran Ur Rahman. 2024. "Is Chatgpt a Menace for Creative Writing Ability? An Experiment." *Journal of Computer Assisted Learning* 40 (2): 919–30. <https://doi.org/10.1111/jcal.12929>.

Noack, Rick. 2024. "As Taliban Limits Options for Afghan Women, Many Lead Secret Lives Online." *The Washington Post*, July 5.

<https://www.washingtonpost.com/world/2024/07/03/afghan-women-girls-online-education/>.

Nordicity. 2025. *Roblox US Economic Impact Assessment 2025*. Roblox. <https://cdn.buttercms.com/kDdTBLOWSamCajkQ2esD>.

Nye, Benjamin D., Dillon Mee, and Mark G Core. 2023. "Generative Large Language Models for Dialog-Based Tutoring: An Early Consideration of Opportunities and Concerns." *CEUR Workshop Proceedings*, July 7, 11. <https://ceur-ws.org/Vol-3487/paper4.pdf>.

O'Donnell, Anna. 2025. "From Crisis to Classroom: How the UN Supports Education in Conflict Zones." United Nations News, July 17. <https://news.un.org/en/story/2025/07/1165417>.

OECD. 2015. *Students, Computers and Learning: Making the Connection*. PISA, OECD. <https://doi.org/10.1787/9789264239555-en>.

OECD. 2021a. *All the Lonely People: Education and Loneliness*. Trends Shaping Education Spotlights, No. 23, OECD. <https://doi.org/10.1787/23ac0e25-en>.

OECD. 2021b. *Measuring Well-Being in the Digital Age*. OECD Going Digital Toolkit Notes, No. 6, OECD. <https://doi.org/10.1787/1891bb63-en>.

Ofcom (Office of Communications). 2024. "Quick Guide to Implementing Highly Effective Age Assurance." [www.ofcom.org.uk](https://www.ofcom.org.uk), February 23. <https://www.ofcom.org.uk/online-safety/illegal-and-harmful-content/age-assurance>.

Office of the Attorney General of California. 2024. "Attorney General Bonta, L.A. City Attorney Feldstein Soto, Announce \$500,000 Settlement with Tilting Point Media for Illegally Collecting and Sharing Children's Data." State of California Department of Justice Office of the Attorney General, June 18. <https://oag.ca.gov/news/>

[press-releases/attorney-general-bonta-la-city-attorney-feldstein-soto-announce-500000](#).

Ohio Legislative Service Commission. 2025. "H.B. 200 136th General Assembly Bill Analysis." Office of Research and Drafting, June 11. <https://www.legislature.ohio.gov/download?key=25563>.

Okolo, Chinasa T., and Marie Tano. 2024. "Closing the Gap: A Call for More Inclusive Language Technologies." Brookings Institution, December 12. <https://www.brookings.edu/articles/closing-the-gap-a-call-for-more-inclusive-language-technologies/>.

Open Source Initiative. 2025. "Open Weights: Not Quite What You've Been Told." *Open Source Initiative*, September 13. <https://opensource.org/ai/open-weights/>.

OpenAI. 2025a. "Enterprise Privacy at OpenAI." OpenAI, September 15. <https://openai.com/enterprise-privacy/>.

OpenAI. 2025b. "Gpt-Oss-120b & Gpt-Oss-20b Model Card." OpenAI, August 5. <https://openai.com/index/gpt-oss-model-card/>.

OpenAI. 2025c. "Introducing Parental Controls." OpenAI, September 16. <https://openai.com/index/introducing-parental-controls/>.

OpenAI. 2025d. "Introducing Study Mode." OpenAI, September 30. <https://openai.com/index/chatgpt-study-mode/>.

OpenAI. 2025e. "Pricing." ChatGPT. <https://chatgpt.com/pricing/>.

OpenAI. 2025f. "Privacy Policy." Open AI, June 27. <https://openai.com/policies/row-privacy-policy/>.

OpenAI. 2025g. "Why Language Models Hallucinate." OpenAI, September 5. <https://openai.com/index/why-language-models-hallucinate/>.

OpenAI. n.d. "Data Controls in the OpenAI Platform." OpenAI. Accessed November 21, 2025. <https://platform.openai.com/docs/guides/your-data>.

OpenAI, Josh Achiam, Steven Adler, et al. 2024. "GPT-4 Technical Report." Preprint, arXiv:2303.08774. March 4, 2024. <https://doi.org/10.48550/arXiv.2303.08774>.

OpenRouter. 2025. "OpenAI." OpenRouter, September 14. <https://openrouter.ai/provider/openai>.

O'Rourke, Meghan. 2025. "I Teach Creative Writing. This Is What A.I. Is Doing to Students." *The New York Times*, July 18. <https://www.nytimes.com/2025/07/18/opinion/ai-chatgpt-school.html>.

Osiurak, François, Jordan Navarro, Emanuelle Reynaud, and Gauthier Thomas. 2018. "Tools Don't—and Won't—Make the Man: A Cognitive Look at the Future." *Journal of Experimental Psychology: General* 147 (5): 782–88. <https://doi.org/10.1037/xge0000432>.

Our World in Data. 2025. "Share of Total Public Education Spending Allocated to Staff Compensation." Our World in Data, May 1. <https://ourworldindata.org/grapher/staff-compensation-as-share-of-total-expenditure-in-public-education-all-levels>.

Ovsyannikova, Dariya, Victoria Oldemburgo de Mello, and Michael Inzlicht. 2025. "Third-Party Evaluators Perceive AI as More Compassionate than Expert Humans." *Communications Psychology* 3 (1): 4. <https://doi.org/10.1038/s44271-024-00182-6>.

Pandit, Rupali, Ina Shastri, and Ajay Surana. 2025. "Design Thinking Approach to Developing Empathetic, Reflective and Inclusive Learners in K-12 Education: A Case Study." *Scope*, 15 (02): 323–351. <https://scope-journal.com/assets/uploads/doc/adfd7-323-351.280047.pdf>.

- Park, Jungho. 2023. "A Case Study on Enhancing the Expertise of Artificial Intelligence Education for Pre-Service Teachers." Preprints, No. 2023052006. May 29. <https://doi.org/10.20944/preprints202305.2006.v1>.
- Park, Peter S., Simon Goldstein, Aidan O'Gara, Michael Chen, and Dan Hendrycks. 2024. "AI Deception: A Survey of Examples, Risks, and Potential Solutions." *Patterns* 5 (5): 100988. <https://doi.org/10.1016/j.patter.2024.100988>.
- Patrinos, Harry, and Nobuyuki Tanaka. 2024. "Education: Innovative Financing in Developing Countries." *Education Working Paper No.1* (April). World Bank. <https://doi.org/10.1596/41476>.
- Patton, Michael Quinn. 2008. *Utilization-Focused Evaluation*. 4th ed. Sage Publications.
- Patwardhan, Tejal, Rachel Dias, Elizabeth Proehl, et al. 2025. "GDPval: Evaluating AI Model Performance on Real-World Economically Valuable Tasks." Preprint, arXiv:2510.04374, October 5, 2025. <https://doi.org/10.48550/arXiv.2510.04374>.
- Pava, Juan N., Caroline Meinhardt, Haifa Badi Uz Zaman, et al. 2025. *Mind the (Language) Gap: Mapping the Challenges of LLM Development in Low-Resource Language Contexts*. Stanford University Human-Centered Artificial Intelligence White Paper. <https://hai.stanford.edu/assets/files/hai-taf-pretoria-white-paper-mind-the-language-gap.pdf>.
- PBS: Firing Line with Margaret Hoover. 2025. "Fei-Fei Li." PBS, May 23. <https://www.pbs.org/wnet/firing-line/video/fei-fei-li-b2ic3j/>.
- Pearson. 2025. *AI Study Tool: Student Outcomes Report: Insights from Connections Academy Online High School Students*. Pearson. <https://plc.pearson.com/sites/pearson-corp/files/2025-08/ai-study-tool-report.pdf>.
- Peele, Stanton, and Archie Brodsky. 1975. *Love and Addiction*. Taplinger.
- Pellegrino, James W., and Margaret L. Hilton, eds. 2012. *Education for Life and Work: Developing Transferable Knowledge and Skills in the 21st Century*. National Academies Press.
- Pentina, Iryna, Tyler Hancock, and Tianling Xie. 2023. "Exploring Relationship Development with Social Chatbots: A Mixed-Method Study of Replika." *Computers in Human Behavior* 140, 107600. <https://doi.org/10.1016/j.chb.2022.107600>.
- Perkins, Mike, Leon Furze, Jasper Roe, and Jason MacVaugh. 2024. "The Artificial Intelligence Assessment Scale (AIAS): A Framework for Ethical Integration of Generative AI in Educational Assessment." *Journal of University Teaching and Learning Practice* 21 (06). <https://doi.org/10.53761/q3azde36>.
- Pichai, Sundar. 2025. "Bringing the Best of AI to College Students for Free." Google (blog), August 6. <https://blog.google/products/gemini/google-ai-pro-students-learning/>.
- Picton, Irene, Christina Clark, and Francesca Bonafede. 2025. *Young People's Use of Generative AI to Support Literacy in 2025*. National Literacy Trust. <https://literacytrust.org.uk/research-services/research-reports/young-people-and-teachers-use-of-generative-ai-to-support-literacy-in-2025/>.
- Pollina, Elvira, and Alvisè Armellini. 2024. "Italy Fines OpenAI over ChatGPT Privacy Rules Breach." *Reuters*, December 20. <https://www.reuters.com/technology/italy-fines-openai-15-million-euros-over-privacy-rules-breach-2024-12-20/>.
- Pontes, Danielle. 2024. "What Is BERT (Bidirectional Encoder Representations from Transformers)?" Zilliz, August 30. <https://zilliz.com/learn/what-is-bert>.
- Postman, Neil. 1990. "Informing Ourselves to Death." Speech delivered to the German Informatics Society. Stuttgart, October 11. <https://web.williams.edu/HistSci/curriculum/101/informing.html>.

- Postman, Neil. 1995. *The End of Education: Redefining the Value of School*. Knopf.
- Potkalitsky, Nick. 2024. "Will Artificial Intelligence Hit a Wall?" *Persuasion* (blog), December 16. <https://www.persuasion.community/p/will-artificial-intelligence-hit>.
- Prifti, Kostina, Jessica Morley, Claudio Novelli, and Luciano Floridi. 2024. "Regulation by Design: Features, Practices, Limitations, and Governance Implications." *Minds and Machines* 34 (2). <https://doi.org/10.1007/s11023-024-09675-z>.
- Project Liberty. 2025. "How to Address the Risks of AI Companions." LinkedIn. <https://www.linkedin.com/pulse/how-address-risks-ai-companions-projectliberty-poghe/>.
- Pugh, Allison. 2024. *The Last Human Job: The Work of Connecting in a Disconnected World*. Princeton University Press.
- Qiang, Christine Zhenwei, and He Wang. 2024. "Who on Earth Is Using Generative AI?" *World Bank* (blog), September 11. <https://blogs.worldbank.org/en/digital-development/who-on-earth-is-using-generative-ai->.
- Qin, Lixia, and Daniel H. Bowen. 2019. "The Distributions of Teacher Qualification: A Cross-National Study." *International Journal of Educational Development* 70 (October): 102084. <https://doi.org/10.1016/j.ijedudev.2019.102084>.
- Quezzaire, Pilar. 2025. "A Frank Chat(Bot) with AI: Using AI for Genuine Inquiry." Webinar. International Baccalaureate Exchange Community Events, April 14. [https://events.zoom.us/j/AuGDfyc6bmc1AyMJodDnquLScDrSdgcrdDGBz8tWDzaTFoi\\_DZiG~AinMsnB6oBo45VGt\\_rcHiXmIx3Q9c-yOgO3NK-5KRItFd1aZAYp7goFEjog](https://events.zoom.us/j/AuGDfyc6bmc1AyMJodDnquLScDrSdgcrdDGBz8tWDzaTFoi_DZiG~AinMsnB6oBo45VGt_rcHiXmIx3Q9c-yOgO3NK-5KRItFd1aZAYp7goFEjog).
- Radboud Universiteit. 2025. "NOLAI." National Education Lab AI, September 9. <https://www.ru.nl/en/nolai>.
- Raffoul, Amanda, Zachary J. Ward, Monique Santoso, Jill R. Kavanaugh, and S. Bryn Austin. 2023. "Social Media Platforms Generate Billions of Dollars in Revenue from U.S. Youth: Findings from a Simulated Revenue Model." *PLoS ONE* 18 (12): e0295337. <https://doi.org/10.1371/journal.pone.0295337>.
- Rahimi, Fatema, Abolghasem Sadeghi, and Soo-Mi Choi. 2025. "Generative AI Meets Virtual Reality: A Comprehensive Survey on Applications, Challenges, and Future Direction." *IEEE Access* 13: 94893–909. <https://doi.org/10.1109/ACCESS.2025.3574779>.
- Raine, Fernande, Susan Rivers, and Michelle Bertoli. 2025. *Teen-centered civics for human thriving: A new blueprint for learning and human flourishing*. History Co: Lab. <https://static1.squarespace.com/static/628d83a9cae7f302fa52eb31/t/69038051177ff606434862c1/1761837137513/Teen-Centered-Civics-Final.pdf>.
- Reed, Jon. 2025. "Congress Isn't Stepping Up to Regulate AI. Where Does That Leave Us Now?" CNET, July 22. <https://www.cnet.com/tech/services-and-software/congress-isnt-stepping-up-to-regulate-ai-where-does-that-leave-us-now/>.
- Reeve, Johnmarshall, Richard M. Ryan, Sung Hyeon Cheon, and Lennia Matos. 2022. *Supporting Students' Motivation: Strategies for Success*. Routledge.
- Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/ 46/ EC (General Data Protection Regulation), Pub. L. No. L 119, 1 (2016). <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679>.
- Reich, Justin. 2025. "What Past Education Technology Failures Can Teach Us about the Future

- of AI in Schools." *The Conversation* (blog), October 3. <https://doi.org/10.64628/AAI.vvdntn96m>.
- Replika.ai. 2025. "Replika." Replika.com, October 31. <https://replika.com>.
- Reschly, Amy L., and Sandra L. Christenson, eds. 2022. *Handbook of Research on Student Engagement*. 2nd ed. Springer International Publishing. <https://doi.org/10.1007/978-3-031-07853-8>.
- Rettberg, Jill Walker. 2022. *ChatGPT Is Multilingual but Monocultural, and It's Learning Your Values*. *Jilltxt* (blog), December 6. <https://jilltxt.net/right-now-chatgpt-is-multilingual-but-monocultural-but-its-learning-your-values/>.
- Reuters. 2025. "OpenAI's Weekly Active Users Surpass 400 Million." *Reuters*, February 20. <https://www.reuters.com/technology/artificial-intelligence/openais-weekly-active-users-surpass-400-million-2025-02-20/>.
- Ring, Hannah Reeves, and Amy R. West. 2015. "Teacher Retention in Refugee and Emergency Settings: The State of the Literature." *International Education Journal: Comparative Perspectives* 14 (3): 106–21.
- Risko, Evan F., and Sam J. Gilbert. 2016. "Cognitive Offloading." *Trends in Cognitive Sciences* 20 (9): 676–88. <https://doi.org/10.1016/j.tics.2016.07.002>.
- Rithm Project. 2025. "About 2." The Rithm Project, October 31. <https://www.therithmproject.org/about>.
- Rivers, Susan E., and Michelle C. Bertoli. 2024. "Using Games to Ignite Teens' Civic and Social and Emotional Learning." *Frontiers in Education* 9 (April). <https://doi.org/10.3389/educ.2024.1322721>.
- RiverSide Learning Center. 2025. *Impact Assessment of Riverside Learning Centre's I CAN Foundation Programme*. [https://www.riversidelearningcenter.in/files/School-offerings/Report\\_impactstudy2025.pdf](https://www.riversidelearningcenter.in/files/School-offerings/Report_impactstudy2025.pdf).
- Robb, Michael B, and Supreet Mann. 2025. *Talk, Trust and Trade-Offs: How and Why Teens Use AI Companions*. Common Sense Media. [https://www.common Sense Media.org/sites/default/files/research/report/talk-trust-and-trade-offs\\_2025\\_web.pdf](https://www.common Sense Media.org/sites/default/files/research/report/talk-trust-and-trade-offs_2025_web.pdf).
- Roberts, Tony, and Marjoke Oosterom. 2024. "Digital Authoritarianism: A Systematic Literature Review." *Information Technology for Development* 0 (0): 1–25. <https://doi.org/10.1080/02681102.2024.2425352>.
- Roblox. 2025a. "Roblox Creator Hub." Roblox, September 13. <https://create.roblox.com>.
- Roblox. 2025b. "Roblox Launches New Learning Hub to Engage Students Through Play." Roblox, July 22. <https://ir.roblox.com/news/news-details/2025/Roblox-Launches-New-Learning-Hub-to-Engage-Students-Through-Play/default.aspx>.
- Rocha, Natalie, and Kashmir Hill. 2025. "Character AI to Ban Children under 18 from Using Its Chatbots." *The New York Times*, October 29. <https://www.nytimes.com/2025/10/29/technology/characterai-underage-users.html>.
- Roy, Palak, Helen Poet, Ruth Staunton, Katherine Aston, and David Thomas. 2024. *ChatGPT in Lesson Preparation: A Teacher Choices Trial*. National Foundation for Educational Research. <https://www.nfer.ac.uk/publications/chatgpt-in-lesson-preparation-a-teacher-choices-trial/>.
- Ruiz, Pati, Kelly Mills, Keun-woo Lee, et al. 2024. *AI Literacy: A Framework to Understand, Evaluate, and Use Emerging Technology*. Digital Promise. <https://doi.org/10.51388/20.500.12265/218>.
- Ryan, R. M., and E. L. Deci. 2000. "Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions." *Contemporary Educational Psychology* 25 (1): 54–67. <https://doi.org/10.1006/ceps.1999.1020>.

Safe AI for Children. 2024. "AI Risks to Children." The Safe AI For Children Alliance, October 23. <https://www.safeaiforchildren.org/risks-of-ai-for-children/>.

Safe and Sounds Schools. 2024. "Comprehensive School Safety." Safe and Sound. <https://www.safeandsoundschools.org/comprehensive-school-safety>.

Sartor, Giovanni, and Francesca Lagioia. 2020. *The Impact of the General Data Protection Regulation on Artificial Intelligence*. Publications Office of the European Union. <https://data.europa.eu/doi/10.2861/293>.

Sawyer, R. Keith, ed. 2014. *The Cambridge Handbook of the Learning Sciences*. 2nd ed. Cambridge University Press. <https://doi.org/10.1017/CBO9781139519526>.

Science Direct. 2025. "Social Desirability Bias." Science Direct, September 29. <https://www.sciencedirect.com/topics/psychology/social-desirability-bias>.

Scragg, Sandy. 2025. "AI Academy Launches." United Federation of Teachers, September 5. <https://www.uft.org/news/news-stories/news-stories/ai-academy-launches>.

Sengeh, David, and Rebecca Winthrop. 2022. *Transforming Education Systems: Why, What, and How*. Brookings Institution. <https://www.brookings.edu/articles/transforming-education-systems-why-what-and-how/>.

Shanahan et al. v. IXL Learning, Inc, No. 24-6985 (9th Cir. files Nov. 19, 2024 2025). [https://www.ftc.gov/system/files/ftc\\_gov/pdf/ShanahanAmicusBriefFiled.pdf](https://www.ftc.gov/system/files/ftc_gov/pdf/ShanahanAmicusBriefFiled.pdf).

Shieh, Evan, Faye-Marie Vassel, Cassidy Sugimoto, and Thema Monroe-White. 2024. "Laissez-Faire Harms: Algorithmic Biases in Generative Language Models." Preprint, arXiv:2404.07475, April 16, 2024. <https://doi.org/10.48550/arXiv.2404.07475>.

Silva, Elise. 2025. "University Students Feel 'Anxious, Confused and Distrustful' about AI in the Classroom and among Their Peers." *The Conversation* (blog), July 16. <https://doi.org/10.64628/AAI.ukwpcxkfr>.

Sindermann, Cornelia, Nana Löchner, Rebecca Heinzelmann, Christian Montag, and Roland W. Scholz. 2024. "The Revenue Model of Mainstream Online Social Networks and Potential Alternatives: A Scenario-Based Evaluation by German Adolescents and Adults." *Technology in Society* 77 (June): 102569. <https://doi.org/10.1016/j.tech-soc.2024.102569>.

Singer, Natasha. 2025a. "Big Tech Makes Cal State Its A.I. Training Ground." *The New York Times*, October 26. <https://www.nytimes.com/2025/10/26/technology/cal-state-ai-amazon-openai.html>.

Singer, Natasha. 2025b. "Microsoft Pledges \$4 Billion Toward A.I. Education." *The New York Times*, July 9. <https://www.nytimes.com/2025/07/09/business/microsoft-ai-education.html>.

Singer, Natasha. 2025c. "OpenAI and Microsoft Bankroll New A.I. Training for Teachers." *The New York Times*, July 9. <https://www.nytimes.com/2025/07/08/technology/chatgpt-teachers-openai-microsoft.html>.

Slocum, Natalie. 2016. "What Is Personalized Learning?" *Education Domain Blog*, February 17. <https://aurora-institute.org/blog/what-is-personalized-learning/>.

Snapchat. 2023. "SPS 2023: What's Next for My AI." Snapchat. <https://newsroom.snap.com/sps-2023-whats-next-for-my-ai>.

Snapchat. 2025. "Early Insights on My AI." Snapchat, September 12. <https://newsroom.snap.com/early-insights-on-my-ai>.

Southern Regional Education Board (SREB). 2025. "AI Tool Procurement, Implementation and Evaluation Checklist." *SREB*. <https://www.sreb.org>.

[org/publication/ai-tool-procurement-implementation-and-evaluation-checklist.](#)

Sparrow, Betsy, Jenny Liu, and Daniel M. Wegner. 2011. "Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips." *Science* 333 (6043): 776–78. <https://doi.org/10.1126/science.1207745>.

Spencer, Michael and Wyndo. 2025. "The Loneliness Economy of AI." *AI Supremacy* (blog), May 20. <https://www.ai-supremacy.com/p/the-loneliness-economy-of-ai>.

Srinivasan, Venkat, and Hemavathi Murthy. 2021. "Improving Reading and Comprehension in K-12: Evidence from a Large-Scale AI Technology Intervention in India." *Computers and Education: Artificial Intelligence* 2 (January): 100019. <https://doi.org/10.1016/j.caeai.2021.100019>.

Staff in the Office of Technology. 2024. "AI Companies: Uphold Your Privacy and Confidentiality Commitments." Federal Trade Commission, January 9. <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/01/ai-companies-uphold-your-privacy-confidentiality-commitments>.

Stasenko, Anastasia, Cormac O'Keefe, and Pierre-Carl Langlais. 2025. "Beyond the Hype: Building Equitable and Sustainable AI for Social Impact." Bibliothèques Sans Frontières, White Paper, October 15. <https://www.bibliosansfrontieres.org/wp-content/uploads/2025/04/White-paper-IA-Advocacy-for-frugal-equitable-IA-042025.pdf>.

Statista. 2025. "Most Used Languages Online by Share of Websites 2025." Statista. <https://www.statista.com/statistics/262946/most-common-languages-on-the-internet/>.

Steiner, Elizabeth D., Ashley Woo, and Sy Doan. 2023. *All Work and No Pay—Teachers' Perceptions of Their Pay and Hours Worked: Findings from the 2023 State of the American Teacher Survey*.

Rand. [https://www.rand.org/pubs/research\\_reports/RRA1108-9.html](https://www.rand.org/pubs/research_reports/RRA1108-9.html).

Steiss, Jacob, Tamara Tate, Steve Graham, et al. 2024. "Comparing the Quality of Human and ChatGPT Feedback of Students' Writing." *Learning and Instruction* 91 (June): 101894. <https://doi.org/10.1016/j.learninstruc.2024.101894>.

Stern, Joanna. 2025. "I Recorded Everything I Said for Three Months. AI Has Replaced My Memory." *The Wall Street Journal*, April 30. <https://www.wsj.com/tech/personal-tech/ai-personal-assistant-wearable-tech-impressions-28156b57>.

Su, Yipeng, and Anna Morgan. 2024. *Promoting Rural Financial Well-Being and Inclusion*. Urban Institute. [https://www.urban.org/sites/default/files/2024-05/Promoting\\_Rural\\_Financial\\_Well-Being\\_Inclusion.pdf](https://www.urban.org/sites/default/files/2024-05/Promoting_Rural_Financial_Well-Being_Inclusion.pdf).

Sweller, John. 1988. "Cognitive Load During Problem Solving: Effects on Learning." *Cognitive Science* 12 (2): 257–85. [https://doi.org/10.1207/s15516709cog1202\\_4](https://doi.org/10.1207/s15516709cog1202_4).

Tang, Xin, Katja Upadyaya, Hiroyuki Toyama, Mika Kasanen, and Katariina Salmela-Aro. 2023. "Assessing and Tracking Students' Wellbeing Through an Automated Scoring System: School Day Wellbeing Model." In *AI in Learning: Designing the Future*, edited by Hannele Niemi, Roy D. Pea, and Yu Lu. Springer International Publishing. <https://doi.org/10.1007/978-3-031-09687-7>.

Teach Access. 2025. "Where AI Meets Accessibility: Considerations for Higher Education." Every Learner Everywhere. <https://www.everylearnereverywhere.org/resources/where-ai-meets-accessibility-considerations-for-higher-education/>.

TeachAI. 2025. "TeachAI." TeachAI. <https://www.teachai.org/toolkit>.

Teacher Task Force. 2025. "Home." Teacher Task Force, October 20. <https://teachertaskforce.org/>

[news/teachers-heart-education-recovery-what-does-latestdata-tell-us-about-state-worlds-teachers.](#)

Tech Justice Law Project. 2025. "TJLP, Young People's Alliance, and Encode File FTC Complaint Over AI Chatbot App Replika's Deceptive Practices, Seek FTC Investigation." Tech Justice Law Project, February 1. <https://techjusticelaw.org/2025/01/31/tjlp-young-peoples-alliance-and-encode-file-ftc-complaint-over-ai-chatbot-app-replikas-deceptive-practices-seek-ftc-investigation/>.

Thompson, Derek. 2025. "The End of Thinking." *Derek Thompson* (blog), July 9. <https://www.derek-thompson.org/p/the-end-of-thinking>.

Tiku, Nitasha. 2024. "An AI Companion Suggested He Kill His Parents. Now His Mom Is Suing." *The Washington Post*, December 10. <https://www.washingtonpost.com/technology/2024/12/10/character-ai-lawsuit-teen-kill-parents-texas/>.

Trabelsi, Wess. 2025. "The Good, the Bad, and the Ugly Science of AI in Education." *Wesstrabelsi* (blog), March 7. <https://wesstrabelsi.substack.com/p/the-good-the-bad-and-the-ugly-science>.

Truong, Nguyen, Kai Sun, Siyao Wang, Florian Guitton, and YiKe Guo. 2021. "Privacy Preservation in Federated Learning: An Insightful Survey from the GDPR Perspective." *Computers & Security* 110 (November): 102402. <https://doi.org/10.1016/j.cose.2021.102402>.

Tunyasuvunakool, Kathryn, Jonas Adler, Zachary Wu, et al. 2021. "Highly Accurate Protein Structure Prediction for the Human Proteome." *Nature* 596 (7873): 590–96. <https://doi.org/10.1038/s41586-021-03828-1>.

Turkle, Sherry. 2024. "Who Do We Become When We Talk to Machines?" *An MIT Exploration of Generative AI* (March). <https://doi.org/10.21428/e4baedd9.caa10d84>.

Turner. 2022. "The National Education Lab AI (NOLAI)." Turner, October 17. <https://www.turner.nl/en/education/the-national-education-lab-ai-nolai/>.

Turner, Amy, Meena Kaushik, Mu-Ti Huang, and Srikar Varanasi. 2022. *Calibrating Trust in AI-Assisted Decision Making*. UC Berkeley School of Information. [https://www.ischool.berkeley.edu/sites/default/files/sproject\\_attachments/humanai\\_capstonereport-final.pdf](https://www.ischool.berkeley.edu/sites/default/files/sproject_attachments/humanai_capstonereport-final.pdf).

Twenge, Jean M., Jonathan Haidt, Andrew B. Blake, Cooper McAllister, Hannah Lemon, and Astrid Le Roy. 2021. "Worldwide Increases in Adolescent Loneliness." *Journal of Adolescence* 93 (December): 257–69. <https://doi.org/10.1016/j.adolescence.2021.06.006>.

Tyack, David, and William Tobin. 1994. "The 'Grammar' of Schooling: Why Has It Been So Hard to Change?" *American Educational Research Journal* 31 (3): 453–79. <https://doi.org/10.2307/1163222>.

UNESCO. 2022. *Transforming Education from within: Current Trends in the Status and Development of Teachers; World Teachers' Day 2022*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000383002/PDF/383002eng.pdf.multi>.

UNESCO. 2023. *Global Education Monitoring Report 2023: Technology in Education: A Tool on Whose Terms?* UNESCO. <https://digitallibrary.un.org/record/4020460>.

UNESCO. 2024. *Global Education Monitoring Report 2024/5*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000391758/PDF/391758eng.pdf.multi>.

UNESCO. 2025. *New UNESCO Report Calls for Multilingual Education to Unlock Learning and Inclusion*. UNESCO, February 18. <https://www.unesco.org/en/articles/new-unesco-report-calls-multilingual-education-unlock-learning-and-inclusion>.

UNESCO Institute for Statistics and Education for All Global Monitoring Report. 2014. "Wanted: Trained Teachers to Ensure Every Child's Right to Primary Education." UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000229913>.

UNESCO and International Task Force on Teachers for Education 2030. 2024. *Global Report on Teachers: Addressing Teacher Shortages and Transforming the Profession*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000388832>.

UNICEF Office of Global Insight and Policy. 2021. "AI Guide for Parents 2021." UNICEF. <https://www.unicef.org/innocenti/media/1361/file/UNICEF-Global-Insight-AI%20guide%20for%20parents-2021.pdf>.

UNICEF-UNESCO. 2025. *Niñas, Niños y Adolescentes Conectados, Informe General de Resultados*. UNICEF Argentina. <https://www.unicef.org/argentina/media/24916/file/NI%C3%91AS,%20NI%C3%91OS%20Y%20ADOLESCENTES%20CONECTADOS%20%20KIDS%20ONLINE%20ARGENTINA%20Informe%20de%20resultados.pdf>.

United Nations Development Programme (UNDP). 2025. "Hamburg Declaration on Responsible AI: Global Leaders Commit to Responsible AI for Sustainable Development." UNDP. [https://www.bmz-digital.global/wp-content/uploads/2025/06/250603\\_Hamburg\\_Declaration.pdf](https://www.bmz-digital.global/wp-content/uploads/2025/06/250603_Hamburg_Declaration.pdf).

United Nations Population Division. 2025. "Population by Age and Sex—Broad Age Groups." Data Portal: Population Division, October 15. <https://population.un.org/dataportal/data/indicators/70/locations/4,903,8,12,16,5505,20,24,660,28,32,51,533,935,36,927,1834,40,31,44,48,50,52,112,56,84,204,60,64,68,535,70,72,76,92,96,100,854,108,132,116,120,124,915,136,140,916,1831,5500,148,152,156,344,446,-158,170,174,178,184,188,384,191,192,531,196,203,408,180,208,901,902,948,262,212,214,910,1832,906,923,218,818,222,226,232,233,748,231,908,1829,5502,>

238,234,242,246,250,254,258,266,270,268,276,288,292,300,304,308,5540,5543,5541,5538,5542,5537,5547,5548,5551,5549,5546,5550,5545,5555,5556,5559,5557,5554,5558,5553,312,316,320,831,324,624,328,332,5503,1503,340,348,352,356,360,364,368,372,833,376,380,388,392,832,400,398,404,296,412,414,417,968,418,904,428,941,422,426,934,430,434,438,440,2087,2088,2089,2090,5504,1859,1500,1501,442,450,454,458,462,466,470,584,474,478,480,175,928,484,954,583,911,1517,492,496,499,500,504,508,104,516,520,524,528,540,554,558,562,566,570,1518,807,912,1833,905,924,580,578,909,1835,512,1401,586,585,591,598,600,604,608,616,957,620,630,634,410,498,638,642,643,646,652,654,659,662,663,666,670,882,674,678,682,686,688,690,2093,2091,2092,694,702,534,703,705,961,90,706,710,931,-728,920,913,5501,925,724,144,275,947,729,740,752,756,760,762,764,626,768,772,776,780,788,792,795,796,798,800,804,784,826,834,840,850,1502,858,860,548,862,704,876,914,922,926,732,900,887,894,716/start/1990/end/2025/table/pivotbylocation.

United Nations Sustainable Development Group (UNSDG). 2024. "Assistive Technology Can Revolutionize Development, Learning and Participation: It's Time Children Everywhere Have Access." UNSDG, July 25. <https://unsdg.un.org/latest/stories/assistive-technology-can-revolutionize-development-learning-and-participation-it%E2%80%99s>.

United Nations Sustainable Development Group (UNSDG). 2025. "Decoding Africa's Energy Journey: Three Key Numbers." UNSDG, January 27. <https://unsdg.un.org/latest/stories/decoding-africa%E2%80%99s-energy-journey-three-key-numbers>.

University of California at Davis. 2025. "Peer & AI Review + Reflection (PAIRR)." UC Davis University Writing Program, February 14. <https://writing.ucdavis.edu/pairr>.

U.S. Department of Education. 2025. "FERPA." U.S. Department of Education, November 19. <https://studentprivacy.ed.gov/ferpa>.

Office of the Surgeon General (OSG). 2023. *Our Epidemic of Loneliness and Isolation: The U.S. Surgeon General's Advisory on the Healing Effects of Social Connection and Community*. U.S. Department of Health and Human Services. <https://www.ncbi.nlm.nih.gov/books/NBK595227/>.

UXPin. 2023. "What Is Progressive Disclosure? Show & Hide the Right Information." *Studio by UXPin*, March 13. <https://www.uxpin.com/studio/blog/what-is-progressive-disclosure/>.

Vegas, Emiliana and Rebecca Winthrop. 2020. "2020: A Year of Turmoil but Also Hope in Education." Brookings Institution, December 22. <https://www.brookings.edu/articles/2020-a-year-of-turmoil-but-also-hope-in-education/>.

Veinott, Beth, Gary A. Klein, and Sterling Wiggins. 2010. "Evaluating the Effectiveness of the PreMortem Technique on Plan Confidence." *Proceedings of the 7th International ISCRAM Conference*, May 9. [https://www.researchgate.net/publication/266485127\\_Evaluating\\_the\\_Effectiveness\\_of\\_the\\_PreMortem\\_Technique\\_on\\_Plan\\_Confidence](https://www.researchgate.net/publication/266485127_Evaluating_the_Effectiveness_of_the_PreMortem_Technique_on_Plan_Confidence).

Veltman, Chloe. 2025. "Ahead of the Holidays, Consumer and Child Advocacy Groups Warn against AI Toys." *NPR*, November 20. <https://www.npr.org/2025/11/20/nx-s1-5612689/ai-toys>.

Venditti, Bruno. 2025. "Ranked: Internet Costs by Country in 2025." *Visual Capitalist*, October 2. <https://www.visualcapitalist.com/internet-costs-by-country-in-2025/>.

Vermont Act No. 63 (2025). <https://legislature.vermont.gov/bill/status/2026/S.69>.

Violino, Bob. 2025. "Generative AI like ChatGPT Is at Risk of Creating New Gender Gap at Work." *CNBC*, May 8. <https://www.cnbc.com/2025/05/08/ai-risk-chatgpt-gender-gap-jobs-work.html>.

Vita-Barrull, Nuria, Verónica Estrada-Plana, Jaume March-Llanes, et al. 2023. "Board Game-Based

Intervention to Improve Executive Functions and Academic Skills in Rural Schools: A Randomized Controlled Trial." *Trends in Neuroscience and Education* 33 (December): 100216. <https://doi.org/10.1016/j.tine.2023.100216>.

Voice 21. 2025. "Home." Voice 21, October 26. <https://voice21.org/>.

Wahn, Basil, Laura Schmitz, Frauke Nora Gerster, and Matthias Weiss. 2023. "Offloading under Cognitive Load: Humans Are Willing to Offload Parts of an Attentionally Demanding Task to an Algorithm." *PLoS ONE* 18 (5): e0286102. <https://doi.org/10.1371/journal.pone.0286102>.

Wall Street Journal. 2019. "Under AI's Watchful Eye, China Wants to Raise Smarter Students." *The Wall Street Journal*, September 19. [https://www.wsj.com/video/series/in-depth-features/under-ais-watchful-eye-china-wants-to-raise-smarter-students/C4294BAB-A76B-4569-8D09-32E9F2B62D19?mod=WSJvidtr\\_pos2](https://www.wsj.com/video/series/in-depth-features/under-ais-watchful-eye-china-wants-to-raise-smarter-students/C4294BAB-A76B-4569-8D09-32E9F2B62D19?mod=WSJvidtr_pos2).

Impact Research. 2024. *AI Chatbots in Schools: Findings from a Poll of K-12 Teachers, Students, Parents, and College Undergraduates*. Walton Family Foundation, May. <https://static.waltonfamilyfoundation.org/bf/24/cd3646584af89e7c668c7705a006/deck-impact-analysis-national-schools-tech-tracker-may-2024-1.pdf>.

Walton Family Foundation. 2025. "The AI Dividend: New Survey Shows AI Is Helping Teachers Reclaim Valuable Time." Walton Family Foundation, June 25. <https://www.waltonfamilyfoundation.org/the-ai-dividend-new-survey-shows-ai-is-helping-teachers-reclaim-valuable-time>.

Walton, Gregory M., Mary C. Murphy, Christine Logel, et al. 2023. "Where and with Whom Does a Brief Social-Belonging Intervention Promote Progress in College?" *Science* 380 (6644): 499–505. <https://doi.org/10.1126/science.ade4420>.

- Wan, Zijun, Jiawei Tang, Linghang Cai, Xin Tong, and Can Liu. 2024. "Breaking the Midas Spell: Understanding Progressive Novice-AI Collaboration in Spatial Design." Preprint, arXiv:2410.20124, October 26, 2024. <https://doi.org/10.48550/arXiv.2410.20124>.
- Wang, Chaoran. 2025. "Exploring Students' Generative AI-Assisted Writing Processes: Perceptions and Experiences from Native and Nonnative English Speakers." *Technology, Knowledge and Learning* 30 (3): 1825–46. <https://doi.org/10.1007/s10758-024-09744-3>.
- Wang, Rose E., Ana T. Ribeiro, Carly D. Robinson, Susanna Loeb, and Dora Demszky. 2025. "Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise." Preprint, arXiv:2410.03017, January 26, 2025. <https://doi.org/10.48550/arXiv.2410.03017>.
- Wangdi, Thinley, and Ringphami Shimray. 2025. "AI-Powered ReadTheory as a Self-Access Learning Platform to Enhance EFL Learners' Reading Enjoyment and Comprehension Skills: A Posthumanist Perspective." *Studies in Self-Access Learning Journal* 16 (2). [https://sisaljournal.org/archives/volume-16/jun25/wangdi\\_shimray/](https://sisaljournal.org/archives/volume-16/jun25/wangdi_shimray/).
- Washington Office of Superintendent of Public Instruction. 2024. "Human-Centered AI: Guidance for K-12 Public Schools." Washington Office of Superintendent of Public Instruction, January 18. [https://ospi.k12.wa.us/sites/default/files/2024-08/comprehensive-ai-guidance-accessible-format\\_0.pdf](https://ospi.k12.wa.us/sites/default/files/2024-08/comprehensive-ai-guidance-accessible-format_0.pdf).
- Watkins, Marc. 2025. "Can We Justify College?" *Rhetorica* (blog), October 31. <https://marcwatkins.substack.com/p/can-we-justify-college>.
- Weaver, D'Andre J. 2022. *Delivering on the Promise of Digital Equity*. Digital Promise. <https://doi.org/10.51388/20.500.12265/166>.
- Webb, Norman L. 2002. *Mathematics Depth-of-Knowledge Levels*. <https://www.modelteaching.com/wp-content/uploads/2019/04/DOK-levels.pdf>.
- Weiner, Steven, Robin Lake, and Jessica Rosner. 2024. *AI Is Evolving, but Teacher Prep Is Lagging: A First Look at Teacher Preparation Program Responses to AI*. Center on Reinventing Public Education. <https://crpe.org/ai-is-evolving-but-teacher-prep-is-lagging/>.
- Weissberger, Alan. 2025. "Sources: AI Is Getting Smarter, but Hallucinations Are Getting Worse." IEE Communications Society Technology Blog, May 10. <https://techblog.comsoc.org/2025/05/10/nyt-ai-is-getting-smarter-but-hallucinations-are-getting-worse/>.
- West, Amy, and Hannah Ring. 2015. "Underresourced, Undervalued, and Underutilized: Making the Case for Teachers in Refugee and Emergency Contexts." *The International Education Journal: Comparative Perspectives* 14: 150–64.
- West, Mark. 2023. "An Ed-Tech Tragedy? Educational Technologies and School Closures in the Time of COVID-19." UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000386701>.
- White House Office of Science and Technology Policy. 2022. "Blueprint for an AI Bill of Rights." *The White House*. <https://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights/>.
- Wilkins, Joe. 2025. "OpenAI Usage Plummets in the Summer, When Students Aren't Cheating on Homework." *Futurism*, August 8. <https://futurism.com/openai-use-cheating-homework>.
- Willingham, Daniel. 2021. *Why Don't Students Like School?: A Cognitive Scientist Answers Questions About How the Mind Works and What It Means for the Classroom*. 2nd Ed. Wiley.
- Winsome Marketing Writing Team. 2025. "The \$250 Stratification: How Google AI Ultra Reveals the Coming AI Class Divide." Winsome Marketing,

September 15. <https://winsomemarketing.com/ai-in-marketing/the-250-stratification-how-google-ai-ultra-reveals-the-coming-ai-class-divide>.

Winthrop, Rebecca, and Eileen McGivney. 2015. *Why Wait 100 Years? Bridging the Gap in Global Education*. Brookings Institution. <https://www.brookings.edu/articles/why-wait-100-years-bridging-the-gap-in-global-education/>.

Winthrop, Rebecca, Adam Barton, and Eileen McGivney. 2018. *Leapfrogging Inequality*. Brookings Institution Press.

Winthrop, Rebecca. 2020. *The Need for Civic Education in 21st-Century Schools*. The Brookings Institution. [https://www.brookings.edu/wp-content/uploads/2020/04/BrookingsPolicy2020\\_BigIdeas\\_Winthrop\\_CivicEducation.pdf](https://www.brookings.edu/wp-content/uploads/2020/04/BrookingsPolicy2020_BigIdeas_Winthrop_CivicEducation.pdf).

Winthrop, Rebecca, Adam Barton, Mahsa Ershadi, and Lauren Ziegler. 2021. *Collaborating to Transform and Improve Education Systems: A Playbook for Family-School Engagement*. Brookings Institution. <https://www.brookings.edu/articles/collaborating-to-transform-and-improve-education-systems-a-playbook-for-family-school-engagement/>.

Winthrop, Rebecca, Youssef Shoukry, and David Nitkin. 2025. *The Disengagement Gap: Why Student Engagement Isn't What Parents Expect*. Brookings Institution. <https://www.brookings.edu/articles/the-disengagement-gap/>.

Wolf, Maryanne. 2018. *Reader Come Home: The Reading Brain in a Digital World*. Harper Collins Publishers.

Wood, Stacy. 2020. "Kids and the Screen-Time Black Hole." *Journal of Consumer Research*, December 30. <http://dqo.xzm.temporary.site/adam-alter-irresistable>.

World Bank. 2023. "World Bank Open Data: Free and Open Access to Global Development Data." World Bank Open Data. <https://data.worldbank.org>.

World Bank Group. 2023. "From Connectivity to Services: Digital Transformation in Africa." World Bank, June 27. <https://projects.worldbank.org/en/results/2023/06/27/from-connectivity-to-services-digital-transformation-in-africa>.

World Bank Group. 2025. *Tracking SDG 7—The Energy Progress Report 2025*. World Bank, June 20. <https://www.worldbank.org/en/topic/energy/publication/tracking-sdg-7-the-energy-progress-report-2025>.

World Economic Forum. 2022. *Artificial Intelligence for Children*. World Economic Forum Toolkit, March 2022. [https://www3.weforum.org/docs/WEF\\_Artificial\\_Intelligence\\_for\\_Children\\_2022.pdf](https://www3.weforum.org/docs/WEF_Artificial_Intelligence_for_Children_2022.pdf).

World Health Organization and UNICEF. 2022. "Almost One Billion Children and Adults with Disabilities and Older Persons in Need of Assistive Technology Denied Access, According to New Report." World Health Organization, May 16. <https://www.who.int/news/item/16-05-2022-almost-one-billion-children-and-adults-with-disabilities-and-older-persons-in-need-of-assistive-technology-denied-access--according-to-new-report>.

World Possible. 2025. "World Possible's RACHEL Server." World Possible. <https://worldpossible.org/>.

Worth, Jack, Sarah Lynch, Jude Hillary, Connie Rennie, and Joana Andrade. 2018. *Teacher Workforce Dynamics in England*. National Foundation for Educational Research.

Xiao, Ping, Yuanyuan Chen, and Weining Bao. 2023. "Waiting, Banning, and Embracing: An Empirical Analysis of Adapting Policies for Generative AI in Higher Education." Preprint, arXiv:2305.18617, May 25, 2023. <https://doi.org/10.48550/arXiv.2305.18617>.

Young People's Alliance, Encode, and Tech Justice Law Project (YPA). 2024. "Deceptive and Unfair Marketing and Design Practices on Replika." YPA, Tech Justice Law Project, and Encode, January 28. <https://techjusticelaw.org/wp-content/>

[uploads/2025/01/Complaint-and-Petition-for-Investigation-Re-Replika.pdf](#).

Youth in AI. 2025. "Youth in AI Summit.: AI for a Sustainable Future: Empowering Youth to Drive Global Impact." Youth in AI. <https://youthinai.org/>.

Yu, Yaman, Yiren Liu, Jacky Zhang, Yun Huang, and Yang Wang. 2025. "Understanding Generative AI Risks for Youth: A Taxonomy Based on Empirical Data." Preprint, arXiv:2502.16383, February 25, 2025. <https://doi.org/10.48550/arXiv.2502.16383>.

Yuan, Ziyang, Xiaoliang Cheng, and Yujing Duan. 2024. "Impact of Media Dependence: How Emotional Interactions between Users and Chat Robots Affect Human Socialization?" *Frontiers in Psychology* 15 (August): 1388860. <https://doi.org/10.3389/fpsyg.2024.1388860>.

Zamora, Gigi. 2023. "Gates Foundation Allocates \$30 Million Towards Advancing AI In Africa." *Forbes*, October 10. <https://www.forbes.com/sites/gigizamora/2023/10/10/gates-foundation-allocates-30-million-towards-advancing-ai-in-africa/>.

Zao-Sanders, Marc. 2025. "How People Are Really Using Gen AI in 2025." *Harvard Business Review*, April 9. <https://hbr.org/2025/04/how-people-are-really-using-gen-ai-in-2025>.

Zhao, Sihang, Shoucong Carol Xiong, Bo Pang, Xiaoying Tang, and Pinjia He. 2025. "Let AI Read First: Enhancing Reading Abilities for Individuals with Dyslexia through Artificial Intelligence." *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (New York), CHI EA '25, April 25, 1–16. <https://doi.org/10.1145/3706599.3720113>.

Zhou, Viola. 2025. "AI Is Reshaping Childhood in China." *Rest of World*, October 1. <https://restof-world.org/2025/ai-china-childhood/>.

Zhuang, Xiaowei, Van Vo, Michael A. Moshi, et al. 2025. "Early Detection of Emerging SARS-CoV-2 Variants from Wastewater through Genome

Sequencing and Machine Learning." *Nature Communications* 16 (1): 6272. <https://doi.org/10.1038/s41467-025-61280-5>.

Zimmerman, Barry J. 2011. "Motivational Sources and Outcomes of Self-Regulated Learning and Performance." In *Handbook of Self-Regulation of Learning and Performance*, edited by Barry J. Zimmerman and Dale H. Schunk. Routledge/Taylor & Francis.

# ANNEX A:

## Example AI literacy frameworks

	ALLit Framework	Digital Promise	China's General AI Education Guide	Uruguay's Marco referencial para la enseñanza de la inteligencia artificial (Reference Framework for Teaching AI)	Day of AI
<b>Year Created</b>	2025	2024	2025	2024	2021
<b>Developers</b>	European Commission (EC), Organization for Economic Co-operation and Development (OECD), Code.org, other experts	Digital Promise	China's Ministry of Education	Ceibal (digital education transformation initiative in Uruguay)	MIT RAISE (Responsible AI for Social Empowerment and Education)
<b>Education Level</b>	Primary and secondary education settings	K-12 education (i.e., early childhood to the end of secondary school)	Primary and secondary schools	K-12 education	K-12 education
<b>Definition of AI Literacy</b>	"The technical knowledge, durable skills, and future-ready attitudes required to thrive in a world influenced by AI. It enables learners to engage, create with, manage, and design AI, while critically evaluating its benefits, risks, and ethical implications."	"The knowledge and skills that enable people to critically understand, evaluate, and use AI systems and tools to safely and effectively participate in an increasingly digital world."		Builds off the 2020 framework developed by Long & Magerko  "A set of competencies that allows people to know and critically evaluate AI technologies, use AI tools, communicate and collaborate with AI, and also promote different ways of thinking that potentially allow creating with AI."	"Includes...the appropriate and productive use of AI, as well as...how AI-powered technologies operate and their applications..., drawbacks and possibilities of [AI], and investigations of its societal and ethical impacts...applications of AI and [how to] interact with it naturally...[and] how AI is relevant in students' lives beyond computer science classrooms."

	AILit Framework	Digital Promise	China's General AI Education Guide	Uruguay's Marco referencial para la enseñanza de la inteligencia artificial (Reference Framework for Teaching AI)	Day of AI
<b>Structure</b>	<p>Domains of AI Literacy:</p> <ol style="list-style-type: none"> <li>1) engaging with AI</li> <li>2) creating with AI</li> <li>3) managing AI</li> <li>4) designing AI</li> </ol>	<p>AI Literacy Practices:</p> <ol style="list-style-type: none"> <li>1) algorithmic thinking, abstraction &amp; decomposition</li> <li>2) data analysis &amp; inference</li> <li>3) data privacy &amp; security</li> <li>4) digital communication &amp; expression</li> <li>5) ethics &amp; impact</li> <li>6) information &amp; mis/disinformation</li> </ol>	<p>AI Education Goals:</p> <ol style="list-style-type: none"> <li>1) cognition</li> <li>2) skills</li> <li>3) thinking</li> <li>4) values</li> </ol>	<p>Dimensions:</p> <ol style="list-style-type: none"> <li>1) what is AI?</li> <li>2) computational learning</li> <li>3) knowledge representation</li> <li>4) computational approach with AI</li> <li>5) use and social impact</li> </ol>	<p>Categories of Units:</p> <ol style="list-style-type: none"> <li>1) AI literacy foundations</li> <li>2) AI applications</li> <li>3) going beyond AI literacy (data activism, creativity, programming)</li> </ol>
<b>Relevant Guardrails</b>	<p>This framework complements the EU's 2022 ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning. These guidelines require that AI systems used in education must have human agency and oversight; transparency; diversity, non-discrimination, and fairness; societal and environmental well-being; privacy and data governance; technical robustness and safety; and accountability.</p>	<p>Digital Promise recommends guidelines and policies that promote the responsible use of AI inside and outside school, and urge educators to ensure students' personal data is protected and caregivers are communicated with about AI tool use at school.</p>	<p>China's Ministry of Education suggests gradual progression of AI use, with lower grades exercising less use than higher grades. Primary level students are not allowed to use open-ended content generation tools without supervision. Teachers are prohibited from relying on AI to answer student questions or evaluate students. The MOE also require guardrails to protect personal data and ensure data security with AI tools in schools.</p>	<p>This framework does not promote specific, policy-oriented guardrails, but it does encourage teacher-led implementation of AI Literacy lessons.</p> <p>Outside of this framework, Uruguay has data protection laws and a national AI strategy that focus on mitigating risks, but they do not focus specifically on children or AI in education.</p>	<p>While Day of AI's curriculum itself doesn't have guardrails built-in, the Day of AI organization will assist school leaders and regulators to create AI policies for their schools.</p>

	AILit Framework	Digital Promise	China's General AI Education Guide	Uruguay's Marco referencial para la enseñanza de la inteligencia artificial (Reference Framework for Teaching AI)	Day of AI
<b>Time with Tech</b>	<p><b>Balanced blend of AI and offline learning</b></p> <p>While this framework does include hands-on engagement with AI tools, several of its lessons focus on foundational skills that can be implemented without AI technologies.</p>	<p><b>Heavy emphasis on hands-on AI use</b></p> <p>While the elements of this framework that focus on understanding and evaluating AI may not always require direct engagement with AI tools, the bulk of the competencies and example lessons in this framework do involve hands-on interaction with AI tools.</p>	<p><b>Age-based progression of AI use</b></p> <p>While China requires that students, beginning at age 6, receive at least 8 hours of AI education a year, this guidance indicates that students in lower grades should gain awareness of AI, while students in higher grades should move to using and creating AI through project-based learning.</p>	<p><b>Lean toward hands-on AI use</b></p> <p>While this framework intentionally includes both conceptual and ethical AI literacy with hands-on engagement with AI tools, it promotes "learning by doing," making hands-on use of AI tools a core component of this model.</p>	<p><b>Balanced blend of AI and offline learning</b></p> <p>While this framework does include hands-on engagement with AI tools, several of its lessons focus on foundational skills that can be implemented without AI technologies.</p>
<b>Source</b>	(European Commission and OECD 2025)	(Ruiz et al. 2024)	(Australian Government Department of Education 2025)	(Capdehourat et al. 2024)	(Day of AI 2025)

## KEY THEMES ACROSS ALL FRAMEWORKS

While these six AI Literacy frameworks are all unique and structured differently, there are many things they have in common. All six frameworks:

1. Include a focus on understanding what AI is and how it works
2. Includes lessons on ethics in AI
3. Include hands-on use with AI tools (though the frequency and type of use differs by framework)

More detail on each framework can be found below:

The **AILit Framework**, currently in draft form and eliciting feedback, outlines four domains of AI Literacy: "Engaging with AI," "Creating with AI," "Managing AI," and "Designing AI." Each domain refers to how AI is used:

- “Engaging with AI” involves using AI as a tool to access information
- “Creating with AI” involves collaborating with AI in creative or problem-solving processes
- “Managing AI” refers to using AI (selectively) as a tool to enhance human work
- “Designing AI” involves shaping how AI systems function

Ethical considerations are embedded across all four domains. For example, “Creating with AI” involves considerations around content ownership and attribution, while “Managing AI” involves assessing whether AI’s role aligns with a user’s goals and values. The AILit Framework emphasizes that AI literacy is comprised of knowledge, skills, and attitudes. Each of the framework’s 22 competencies, which extend across the four domains, is meant to grow and develop specific knowledge, skills, and attitudes of learners. The framework provides example primary and secondary lessons for each of its competencies.

**Example lessons in the AILit Framework:**

<p><b>Domain:</b> Creating with AI</p> <p><b>Competency:</b> Collaborate with generative AI systems to elicit feedback, refine results, and reflect on thought processes.</p> <p><b>Knowledge:</b> AI systems gather new data from interactions with users; decisions, processes, and outputs may be directly influenced by inputs in real time.</p> <p><b>Skills:</b> computational thinking, creativity</p> <p><b>Attitudes:</b> innovative, adaptable</p> <p><b>Lesson (secondary grades):</b> Use an AI coding assistant to fix errors and modify code for a video game, then reflect on how the tool affected the debugging process.</p>	<p><b>Domain:</b> Managing AI</p> <p><b>Competency:</b> Decide whether to use AI systems based on the nature of the task.</p> <p><b>Knowledge:</b> (1) AI excels at pattern recognition and automation but lacks emotions, ethical reasoning, context, and originality. (2) Ethical AI design encompasses fairness, transparency, explainability, accountability, respect for privacy, and legal compliance.</p> <p><b>Skills:</b> problem solving, computational thinking</p> <p><b>Attitudes:</b> responsible, innovative</p> <p><b>Lesson (primary grades):</b> Consider everyday tasks (e.g., writing a birthday card) and assess when AI use is appropriate, considering the need for individuality, creativity, or human judgment.</p>
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

The **Digital Promise Framework** defines three modes of engagement for AI literacy: understand, evaluate, and use. These are interconnected, demonstrating that engaging with AI in these three ways together leads to robust AI literacy. To understand AI, this framework suggests instruction on the history of AI, how it works, and how it’s being used. To evaluate AI, this framework emphasizes centering human judgement and justice, concepts that serve as baselines for how and why we should evaluate and be critical consumers of AI tools. To demonstrate students’ ability to understand and evaluate AI, this model describes 6 AI literacy practices:

- Algorithmic thinking, abstraction, & decomposition
- Data analysis & inference
- Data privacy & security
- Digital communication & expression
- Ethics & impact
- Information & mis/disinformation

Digital Promise also outlines three different types of educational AI use: interact, create, and problem-solve. For each use type, they provide examples of how students may apply the literacy practices.

**Example applications of Digital Promise’s AI literacy practices:**

<p><b>Interacting with AI:</b>  <b>Algorithmic thinking, abstraction &amp; decomposition</b>          Understand how AI tools make automated recommendations and how to influence recommendations for specific goals.</p> <p><b>Digital communication &amp; expression</b>          Responsibly engage with and share AI-influenced or AI-created content online.</p>	<p><b>Creating with AI:</b>  <b>Ethics &amp; impact</b>          Consider how your creation aligns to personal values/ethics. Ensure appropriate and ethical application of synthetic content (e.g. citing appropriately, mitigating harm of deep fakes, etc.). Become aware of the impact of consumption of GenAI-created content on self, others, environment, and society.</p> <p><b>Information &amp; mis/disinformation</b>          Identify mis/disinformation in synthesized outputs/products. Evaluate whether you are perpetuating false information when creating (e.g. do not include mis/disinformation in prompts).</p>	<p><b>Problem Solving With AI:</b>  <b>Data analysis &amp; inference</b>          Evaluate the sources, contexts, and organization of data on which the algorithm is trained.</p> <p><b>Data privacy &amp; security</b>          Consider what personal identifiable information (PII) or sensitive data you are collecting in your approach to problem solving and how to protect user privacy and security.</p>
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

**China’s Ministry of Education** issued an *AI Education Guide for Primary and Secondary Schools*. This guide defines four key areas of AI literacy: cognition, skills, thinking, and values. For each of these areas, the framework outside goals for primary, junior high, and senior high students.

## Example Goals

Skills	Values
<p><b>Goal for primary students:</b> Develop basic AI skills by using simple tools, visual programming, and practicing data handling.</p> <p><b>Goal for junior high students:</b> Apply AI to real-world problems through project-based learning and building simple intelligent agents.</p> <p><b>Goal for senior high students:</b> Focus on innovative AI use by building simple models and creating interdisciplinary solutions with intelligent tools.</p>	<p><b>Goal for primary students:</b> Deepen ethical understanding by recognizing AI's strategic role in innovation and evaluating misinformation risks in generative technologies.</p> <p><b>Goal for junior high students:</b> Deepen ethical awareness by recognizing AI's strategic value and evaluating risks like misinformation in generative applications.</p> <p><b>Goal for senior high students:</b> Promote social responsibility by examining AI sovereignty within national tech strategies and balancing innovation with ethical risks.</p>

**Uruguay's Reference Framework for Teaching Artificial Intelligence** aims to build students' AI literacy and computational thinking skills; encourage critical, creative, and ethical engagement with AI; and prepare students to understand, use, and create AI systems responsibly and thoughtfully. The framework outlines six core principles that should underpin all AI learning activities: equity, collaboration, creativity, autonomy, critical perspective, and active methodologies. Pedagogically, the framework is organized into five dimensions:

- What is AI?
- Representation of knowledge
- Computational learning
- Computational approach with AI
- Ethical use and social impact

This model views AI literacy as a multidisciplinary topic, and encourages combining computing, ethics, and social sciences. It also encourages project-based learning, having students learn real examples of AI use and doing hands-on experimentation with AI tools.

## Example learning goals for each AI literacy dimension

<p><b>Representation of knowledge</b></p> <ul style="list-style-type: none"> <li>• Identify sensors and data sources</li> <li>• Understand how data representation influences AI behavior</li> <li>• Critically analyze AI outputs and potential errors</li> </ul>	<p><b>Computational approach with AI</b></p> <ul style="list-style-type: none"> <li>• Identify real-world problems suitable for AI solutions</li> <li>• Apply AI tools for analysis, prediction, or creative tasks</li> <li>• Evaluate AI's strengths and weaknesses</li> </ul>	<p><b>Ethical use and social impact</b></p> <ul style="list-style-type: none"> <li>• Identify positive and negative social impacts of AI</li> <li>• Recognize bias in data and algorithmic decisions</li> <li>• Imagine future applications and discuss their societal implications</li> </ul>
--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

The **Day of AI** curriculum was developed at MIT in 2021 to prepare K–12 students to be successful, responsible, and engaged in an AI world. This curriculum provides free resources full of hands-on AI literacy lessons accompanied by professional development for teachers. The curriculum has been taught in all 50 U.S. states and over 110 countries. For example, in August 2025, Rwanda's Ministry of Education and the Rwandan Education Board hosted a four-day national training for 150 master teachers on the Day of AI curriculum. This curriculum is organized by three themes: AI literacy foundations, AI applications, and going beyond AI literacy. For each theme, the curriculum provides units for students aged 4–7, 8–10, 11–13, and 14+.

## Example Units in the Day of AI Curriculum

<p><b>AI literacy foundations</b></p> <ul style="list-style-type: none"> <li>• AI for early elementary (ages 4–7)</li> <li>• What is AI? (ages 8+)</li> <li>• How do machines learn? (ages 8+)</li> <li>• How do machines create? (ages 11+)</li> <li>• AI and ethics materials (ages 11+)</li> </ul>	<p><b>AI applications</b></p> <ul style="list-style-type: none"> <li>• Using AI for creativity (ages 4–7)</li> <li>• AI and the creative arts (ages 8+)</li> <li>• ChatGPT in schools (ages 11+)</li> <li>• How do AI tools affect our rights? (ages 11+)</li> <li>• AI legislation in schools: empowering young educators (ages 11+)</li> <li>• Understanding AI in social media (ages 11+)</li> <li>• The brain behind the bot: chatbots, bias, and productive uses of AI for kids (ages 12+)</li> </ul>	<p><b>Going beyond AI literacy</b></p> <ul style="list-style-type: none"> <li>• AI programming courses with blocks (ages 8+)</li> <li>• How are we quantified by AI? (ages 14+)</li> <li>• AI programming courses with Python (ages 14+)</li> <li>• Careers in AI (ages 14+)</li> <li>• Making sense of our surroundings (ages 14+)</li> </ul>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

# ANNEX B:

## Acknowledgements

We extend our deepest gratitude to the global community whose engagement made this study possible. Teachers, students, parents, education and technology experts, leaders, and others from around the world generously shared their time, insights, and experiences, enriching our work with diverse perspectives.

### STEERING COMMITTEE

**Armand Doucet**

*Former Senior Advisor for Artificial Intelligence in Education, Government of New Brunswick, Canada*

**Armando Estrada Zubía**

*Cofounder and Executive Director, Via Educación, Mexico*

**Bassem Sa'ad**

*Chief Executive Officer, Queen Raina Foundation for Education and Development; Former European Relations Senior Officer, Ministry of Planning and International Cooperation, Jordan*

**Benjamin Piper**

*Director of Global Education Program, Gates Foundation*

**Carolina Piñeros**

*Executive Director, Red PaPaz, Colombia*

**Chernor Bah**

*Minister of Information and Civic Education, Government of Sierra Leone*

**Claudia Costin**

*President, Salto Institute; Former Secretary of Education, Rio de Janeiro City Hall, Brazil*

**Cristóbal Cobo**

*Senior Education and Technology Policy Expert, World Bank*

**David Edwards**

*General Secretary, Education International*

**Denis Mizne**

*Chief Executive Officer, Lemann Foundation, Brazil*

**Femi Longe**

*Global Bitcoin Lead, Human Rights Foundation*

**Fengchun Miao**

*Former Chief of the Unit for Technology and AI in Education, UNESCO; Former Director, National Center for ICT in Education, Ministry of Education, China*

**Frank van Cappelle**

*Head of the Office of Innovation, UNICEF*

**George Papandreou**

*General Rapporteur on Democracy, Council of Europe Parliamentary Assembly; Former Prime Minister of Greece*

**Glenn M. Kleiman**

*Senior Advisor, Stanford Accelerator for Learning, U.S.*

**Isabelle Hau**

*Executive Director, Stanford Accelerator for Learning, U.S.*

**Jean-Claude Brizard**

*President and CEO, Digital Promise, U.S.*

**Jim Knight**

*Member, House of Lords; Former Minister of State for Schools, UK*

**Jon Valant**

Senior Fellow and Director of the Brown Center on Education Policy, Brookings Institution, U.S.

**Jordan Shapiro**

Associate Professor, Temple University, U.S.

**Juan-Pablo Giraldo**

Education and Innovation Specialist, UNICEF

**Leandro Folgar Ruétalo**

Vice President of Innovation, Universidad Católica del Uruguay; Former President, Fundación Ceibal, Uruguay

**Michael Trucano**

Non-Resident Fellow, Brookings Institution, U.S.

**Olli-Pekka Heinonen**

Director General, International Baccalaureate; Former Director General of the Finnish National Agency for Education, Finland

**Pat Yongpradit**

Chief Academic Officer, Code.org; Lead, TeachAI, U.S.

**Pati Ruiz**

Director of Learning Technology Research, Digital Promise, U.S.

**Punya Mishra**

Professor and Director of Innovative Learning Futures at the Learning Engineering Institute, Mary Lou Fulton College of Teaching and Learning Innovation, Arizona State University, U.S.

**Sara Ruto**

Program Officer, Echidna Giving; Former Chief Administrative Secretary, Ministry of Education, Kenya

**Shabana Basij-Rasikh**

Cofounder and President, School of Leadership, Afghanistan

**Shafika Isaacs**

Chief of Section, Technology and AI in Education, UNESCO

**Shaveta Sharma-Kukreja**

Chief Executive Officer and Managing Director, Central Square Foundation, India

**Stéphan Vincent-Lancrin**

Deputy Head of Division and Senior Economist and Analyst, OECD

**Urvashi Sahni**

Founder, President, and Chief Executive Officer, Study Hall Educational Foundation, India

**Verna Lalbeharie**

Executive Director, EdTech Hub; Former Director of Digital Teaching and Learning for the North Carolina Department of Public Instruction, U.S.

## STUDY PARTICIPANTS

We are incredibly grateful to our resource experts, Delphi panelists, focus group participants, consultation participants, and interview participants, who dedicated their time and shared their experience with our team. The following list includes the study participants who gave us permission to recognize their contributions. We are deeply grateful to all other study participants not listed, including the students we interviewed, none of whom are listed here to protect their identity.

Abby David	Barbara Treacy	Cosmas Kamwendo	Elisabeth O’Byron
Adesina Oluwatosin Israel	Barry Cottam-Howarth	Craig Findlay	Elizangela de Nazaré Cardoso Damasceno
Adewale Orebanwo	Billy Baum	Cristine Legare	Emmanuel Larbi
Adjei Claudia	Bonginkosi Ntosome	Dalila Maite Rosa Sena	Enrique Lemitton
Adriano de Sousa Sá	Brian Adams	Dana Suskind	Eric Anctil
Alan Al Ahmad	Carla Aerts	Daniel Barcay	Eric Hudson
Álex Gerardo Hernández Silva	Carolina Augustina Vasquez de Estraver	Daniela Cardoso da Silva	Felipe Andrés Contreras Bravo
Alireza Alidoust	Carol Yu	Danny Martinez	Felix Dzanku
Amina Yekhlef	Celine Perea	Dave Smith	Francesca Mazzieri
Amokelo Maweya	Charles Fadel	David Castañeda	Francilda Fonseca
Ana María Raad	Chernor Bah	David Hay	Gandonu Temitope
Anitha A	Chiamaka Martha Nkwuenu	David Monahan	George Mutale
Anjum Ara	Christin Monroe	David Montagu	Giancarlo Brotto
Ann Obieje	Chris Ludwig	Delphine Le Serre	Gita Luz
Ann O’Keefe	Clarissa Paz de Menezes	Dinu Raheja	Glenn M. Kleiman
Armand Doucet	Cloé Ward	Drew Edwards	Gouri Gupta
Aslı Karabenli	Connor Akers	Dyone Souza de Andrade	Gurpreet Kaur Arora
Audaci Maria de Lima Silva	Corrine Cross	Edison Muyala	Hakeoung Hannah Lee
		Elham Tabassi	

Harsha Makan	Kavya M	Marjana Skenduli	Natalie Johnson
Henré Benson	Kevin Martin	Matt McGuire	Natasha Armstrong
Henry Kofi Rockson Junior	Khuteja Bi	Matt Pulley	Nazareen Ebrahim
Heo Young Joo	Kim Min Joo	Matt Strikwerda	Nazia
Imogen Casebourne	Lance Eaton	Matthew Jessop	Neha Yadav
Isabel Anabela Vergara Angulo	Larry Molinaro	Mehar Sultana	Nelson Lasso Chiro
Jacqueline da Silva Santos Batista	Lauren Paer	Micah Stickler	Neuza Sofia Guerreiro Pedro
Jade Ackers	Lawal Ibironke Remilekun	Micheal Bamidele	Nidhi Singh
James Basham	Lawrell Munsami	Michael Harvey	Nkosilathi Dube
Jara Medina	Léonie Aubut	Michael McNabb	Nomsa Daka
Jen Scott	Letícia Matia de Souza Cortes	Michael Serazio	Nupke Sunu Xorlanya Akosua
Jeremy Goldstein	Lindsay Dombrowski	Michael Sullivan	Ogunmoroti Ayodeji O.
Jim Knight	Linice Sanga	Michael Trucano	Omashani Naidoo
Jody Britten	Luke Stannard	Miloš Jeremić	Osondu Uchechukwu L
Jordan Wallace-Wolf	Madinkgoana Mirriam Msiza	Mogomotsi Molefi	Otodei Nii Odoi Joseph
Joseph Smith	Manuel Arthur Gil Jara	Mokgadi Obakeng	Park Hye Ran
Joy Amarachi Ifeanyi	Monica Gagnon	Mónica Genoveva Lesama Matute	Pati Ruiz
Julian Fraillon	Margaret LaRaia	Monique Arsenault	Patience Ogechi Agonsi
Justin Kirby	María Carmen Romero Monzón	Nadia Abdallah	Patricia Jacoba Sáñez López Torres
Karen Murphy	María Claudia Uribe Salazar	Nadine Leblanc	Peter Rafferty
Kathleen Doyle	Marianela Alejandra Vergaalera Cortés	Najlah Al Bin Ali	Philip McRae
Kathy Hirsh-Pasek	Marie-Laure Pokou	Natalia Edisherashvili	Pilar Quezzaire

Poonam Dixit	Robin Street	Sarika Sharma	Stephen Manchisi
Poppy Ngidi	Rod Scott	Seidu El-Ham	Suman Singh
Quintres Sefala	Roshaan Suren	Sergio Lance Somerville	Sumithra
Rachel Romeo	Ross Hall	Shamsiya	Sydney Munsterman
Rachel Urquhart	Ryan Layton	Sheeba A	Syeda Amina
Radhika Yadav	Sabba Quidwai	Shirley King	Thomas Thompson
Rafael Parente	Sacha van Straten	Simon Collier	Tracey Bryan
Rafael Rebelo	Sadia Zia	Simone Natale	Vanessa Lusa
Ragini Upadhyay	Saeed Ullah Khan	Simran Mulchandani	Vincent Sakwata
Rapelang Rabana	Sam Scott	Sizwe Zulu	Vriti Saraf
Richard Kingleza	Samantha Skelton	Sonu M	Wayan Vota
Rida Karim	Sara Ratner	Sonya Hull	Yishay Mor
Rivaldo Bevenuto de Oliveira Neto	Sara Ruto	Sophie Elliott	

## 17 ROOMS INITIATIVE

We are also grateful to the participants of Room 4 (Quality Education) in the 17 Rooms initiative convened by Brookings and the Rockefeller Foundation. Their insights and discussion helped us advance our thinking on generative AI and students' experiences. We offer deep gratitude to those listed below, whose diverse contributions helped make this report possible.

Aria Finger	Drew Bent	Kershlin Krishna	Stephen Jull
Ben Gomes	Ellen Pack	Lilach Mollick	Vibhu Mittal
Dan Carroll	Jacob Taylor	Pat Yongpradit	
Denis Mizne	James Slavet	Philip McRae	

## OUR TEAM AND COLLABORATORS

We offer deep gratitude to those listed below. Their diverse contributions helped make this report possible.

Ashleigh Ekwenugo	Humyra Karim	Linice Sanga	Park Kiwoong
Annick Arsenault-Carter	Izzy Taylor	Marian Licheri Hougaard	Pooja Sharma
Chinenye Agwuncha	Jayant Gauri	Mashhood Bhat	Punya Mishra
Dayse Oliveira	Jennifer O'Donoghue	Michael Maffie	Samarth Shukla
Elham Tabassi	Joel Mitchell	Nina Ng	Steve McDonald
Emily Morris	John M. Costello	Nova Escola (institution)	Surabhi Yelsangikar
Gaia Bernstein	Joysy John MBE	Ogulkeyik Hudayberdiyeva	Valeria Duarte
Hans Cabra	Lawrencia Owusu		

# ANNEX C:

## Research methods

This research is grounded in the concept of a premortem. A premortem is a strategic planning exercise that anticipates potential failures before implementing an innovation. Unlike a **postmortem**, which examines failure after occurrence, a premortem assumes failure has already occurred and works backward to identify the causes of failure (Burns and Winthrop 2025). Premortems have two main benefits: first, they reduce overconfidence, which emerges when teams with shared goals construct favorable narratives about their initiatives, creating a self-reinforcing cycle where contradictory information is dismissed and confidence becomes increasingly divorced from reality (Veinott et al. 2010). Secondly, premortems also eliminate social desirability bias by creating structured opportunities for honest critique, overcoming the natural reluctance to express skepticism about promised benefits (Science Direct 2025). This dual function generates higher-quality information and more accurate risk assessments, ultimately enhancing innovation outcomes through more realistic evaluation processes (Klein 2021).

### RESEARCH QUESTIONS

- What are the potential negative risks that generative AI poses to the education of children and youth?
- Assuming these potential risks, what can we begin to do now to prevent them and maximize the potential benefits of AI?

### METHODOLOGY

Our primary methods for investigating these research questions included: conducting a literature review, conducting qualitative research via focus groups and group interviews, and verifying our data through a Delphi Panel and education stakeholder consultations.

### LITERATURE REVIEW

Guided by the two research questions and to better situate study participants' experiences within a theoretical context, we conducted an extensive review of studies on AI in education published after November 2022. Given the nascent state of this field, the rapid pace of AI development, and the scarcity of experimental and longitudinal research, the review adopted a pragmatic approach, also incorporating preliminary findings and smaller-scale investigations, policy documents, essays, and case studies. Researchers also accessed news reports on AI in education, AI thought pieces from other disciplines, and essays by respected AI in education leaders. In total, we reviewed over 425 articles, using these insights to inform our analysis and findings.

### INITIAL DATA COLLECTION

**Focus groups (N = 72):** This was conducted primarily with education experts—those working in the field of AI in education, university researchers, and those associated with governments, businesses, and NGOs. Within this report, they are referred to as “experts.”

Focus groups were conducted over two sessions, with session one focused on the first research question (potential negative risks of AI) and session two focused on addressing these risks and harnessing benefits.

**Group interviews (N = 226):** Interviews were conducted with 117 teachers, 68 students, and 41 parents. Interview questions focused on four questions: AI's uses in education, its benefits, its risks, and what to do about these risks. Interviews were one hour in length.

All data were coded and analyzed, and themes were developed using qualitative research software. Quotes from experts, teachers, parents, and students can be found throughout this report.

## VERIFYING THE DATA

**Delphi panel (N = 21):** For questions that appeared to have no definitive answers—for example, questions around AI adoption, promoting agency, and rethinking the purpose of education in an AI world, we deployed a Delphi panel, a structured, anonymous, iterative approach that harnesses the collective wisdom of a group of experts. The observations of the Delphi panel are integrated throughout this report, particularly in Section VI.

**Consultations (N = 232):** We conducted consultations with a wide variety of stakeholders to inform each stage of the research process: getting feedback on the methodology, stress testing data collection instruments, surfacing gaps in preliminary findings, and final data analysis. The participants in these consultations ranged from high-level decisionmakers (including those on the task force steering committee) to funders and practitioners such as ed-tech leaders, nongovernmental organization representatives, child development specialists, learning scientists, and technologists. Included in these consultations were four task force meetings with steering committee members who engaged in guidance and feedback on the study from September 2024 to October 2025.

In total, our team engaged with 505 unique participants across the initial data collection and data verification phases. Some individuals took part in both phases, which is why the combined count exceeds 505.

## COUNTRIES

The study was global in scope, and we drew on a wide range of literature available, primarily in English. Study participants—those participating in either focus groups, group interviews, the Delphi panel, or consultations—were from the 50 countries listed below:

1. Afghanistan	14. France	27. Malaysia	40. South Africa
2. Albania	15. Germany	28. Mexico	41. South Korea
3. Australia	16. Ghana	29. Netherlands	42. Spain
4. Brazil	17. Greece	30. Nigeria	43. Taiwan
5. Canada	18. Hong Kong	31. Pakistan	44. Thailand
6. Chile	19. Indonesia	32. Paraguay	45. Tunisia
7. China	20. India	33. Peru	46. United Arab Emirates
8. Colombia	21. Israel	34. Portugal	47. United Kingdom
9. Denmark	22. Italy	35. Romania	48. United States
10. Ecuador	23. Jordan	36. Saudi Arabia	49. Uruguay
11. Egypt	24. Kenya	37. Serbia	50. Zambia
12. Estonia	25. Kingdom of Bahrain	38. Sierra Leone	
13. Finland	26. Malawi	39. Singapore	

BROOKINGS

*©The Brookings Institution, January 2026. All Rights Reserved.*