MEMO *Published November 17, 2026*

Guardrails, Not Gates: Empowering the AI Generation



By: **Maya Schmidt**

Today's kids will be the first generation to grow up with AI woven into nearly every part of their lives. It could become a tutor that helps them learn, a tool that gives kids with disabilities greater independence, or a creative engine that helps them solve problems

we haven't imagined yet. And learning to use AI will increasingly be part of preparing kids for the world they'll inherit.

But growing up with AI also means growing up with risks we are only beginning to understand. These products and tools can collect deeply personal information, be designed to keep kids hooked, and expose them to harmful or dangerous content. Although AI is different from social media, we still know far too little about what sustained interaction with AI could mean for children's social, emotional, and cognitive development. But there are also lessons from the social media safety debate that can translate into the AI ecosystem.

In this memo, we lay out a path to protect kids from AI's biggest risks without shutting them out of its benefits. The answer isn't to put up gates that keep kids away from AI. It's to build guardrails that make AI safer for them to use and arm parents and kids with the resources they need to navigate it. Companies that put AI-enabled products and tools in children's hands should be responsible for making them safe from the start. Parents should have meaningful tools to protect their kids. And policymakers should invest in understanding the risks we don't yet fully understand. Done right, we can give kids the freedom to explore what AI makes possible without asking them to bear its biggest risks.

How Kids Use AI

When people think about kids using AI, chatbots or companion toys likely come to mind first. Those uses are important, but AI is showing up in many other parts of kids' lives. Understanding all the different ways kids encounter AI is important because each has its own benefits and risks. The following are some, but not all, of the ways kids use AI:

- **Assistive Technology**: These tools help children with disabilities communicate, learn, and navigate daily tasks. Under federal law, assistive technology is defined as "any item, piece of equipment, or product system... that is used to increase, maintain, or improve the functional capabilities of a child with a disability."[1]
- **AI Companion Toys**: Physical toys designed and marketed as friends or pets, embedded with AI that allows children to engage in real-time, two-way spoken

conversations.

- **Educational AI**: Platforms and tools that are designed specifically for learning and can adapt lessons and support based on a child's real-time input.
- **Chatbots**: General-purpose software that uses AI to simulate human conversation. Designed to adapt to user input, chatbots can respond differently to different users and change their behavior over time.
    - This category can include **companion chatbots**, which are primarily designed to simulate interpersonal relationships with users
    - It can also include **general-purposechatbots**, which are designed for open-ended conversation across a broad range of topics. These chatbots may simulate interpersonal relationships, but they can also provide information, facilitate productivity, and support creative projects.
    - Finally, this category also includes **task-specific chatbots**, which are designed for and restricted to a specific function, such as customer service or technical support.
- **Voice Assistants**: Home devices that use AI to respond to spoken commands, answer questions, perform tasks, and control connected environments.

What Kids Can Gain

AI isn't just a new risk for kids—it's also a powerful new tool. It can help children with disabilities communicate, give students individualized support, and put powerful tools in the hands of young people solving problems in their own communities. And as AI reshapes school, work, and everyday life, kids will need to learn how to use it well. We're already seeing what that potential can look like:

Access & Independence

AI can make assistive technology more useful and give children with disabilities greater independence. While the technology is relatively new, developers are finding ways to

make these tools more helpful. For example, AI has enabled tap-to-speak devices to better reflect natural language patterns, conveying a wider range of emotions and intentions.[2] It has also been used to help people with dyslexia process large amounts of reading material and people with executive functioning challenges organize their thoughts.[3] AI can give kids and teens greater independence and make it easier to navigate a world that wasn't always designed for them.

Personalized Learning

AI can also give students more individualized support and help teachers use their time more effectively. Educational AI tools are designed to help students work through problems without simply giving them the answers. One example is Khanmigo, an AI assistant designed for tutoring. One teacher said the tool allows students who can work at their own pace to learn independently while he focuses on those who need more help.[4] MATHia is AI-powered math software designed to teach students not only what they got wrong but why.[5] One study found that students saw substantial improvement in math testing performance after using the program.[6] "AI can make measurable contributions to learning, particularly in districts with limited resources. AI can also expand access to personalized services that would ordinarily be too expensive.

Creativity & Innovation

AI can open doors for kids and teens to develop creative solutions to problems in their communities. For example, a 17-year-old set out to protect a river two miles from his home after learning it was threatened by excessive water withdrawals and flooding. The problem? There were only two gauges measuring water levels, and the Ipswich River Watershed Association couldn't afford to install another. Rather than collect more data, he applied his AI and machine-learning skills to develop project IRIS, which enables real-time flood prediction.[7]

Separately, a high school senior created SensoRy AI: a solar-powered sensor network that uses AI to detect early ignitions that could lead to wildfires.[8] Another teenager was a finalist in a Society for Science competition, where she equipped an unmanned aerial vehicle with AI to detect disease in crops.[9] These kids aren't just learning how to use AI —they're using it to solve real problems.

Digital Literacy

Kids are going to navigate an AI-enabled world for the rest of their lives. Early exposure can help them build the skills they'll need— not only to survive but also to thrive in it. That means learning not just how to use AI, but how to use it responsibly, recognize its limitations, and know when not to trust it. The key is to establish guardrails that curb the most dangerous risks of the technology while still allowing kids to develop the literacy skills and responsible relationships with AI they will need in the future.

Where Kids Face Risks

The risks AI poses to kids are as real as its benefits. Children are more trusting, more easily influenced, and still developing the judgment to recognize when something is wrong. That makes them more vulnerable when AI products fail them: their personal information can be collected and exposed, products can be designed to keep them hooked, and they can encounter content ranging from inaccurate to dangerous. And because the technology is so new, some of the most important questions—including how AI shapes a child's development over time—remain unanswered.

Data Privacy

Kids can share an enormous amount of personal information with AI—and they may have no idea where that information goes, how it's protected, or who it's shared with. In January, ethical hackers Joseph Thacker and Joel Margolis discovered that the maker of AI companion toy Bondu inadvertently left more than 50,000 transcripts of children's conversations accessible to anyone with a Gmail account.[10] Once a child's personal information is leaked, it's hard to undo that damage. And it's not the first time children's data has been exposed.[11]

Leaks of personal information can put kids at risk of identity theft, financial fraud, and even stalking and harassment. But the risks go beyond obvious leaks of sensitive information. Even non-sensitive data can reveal deeply personal things about a child.[12] The widespread collection, retention, and sale of kids' personal information allow companies to track and profile them, and those profiles can follow them throughout their lives, even when they're accurate. Most kids don't understand those

long-term risks, and most parents do not have the time to decipher every company's privacy policies to track how their kids' data is being used. Plus, as AI gets better at spotting patterns, ordinary information can reveal more than kids or their parents ever intended. Even data that seems harmless today can adversely affect a child years from now.

Design-Related Risks

Developmental and child psychologists are concerned that kids could form unhealthy attachments to AI. The same features that make these products appealing to children can also make them unsafe. Experts warn that when AI is optimized for engagement, as social media platforms are, it can be remarkably effective at making a child feel like they need it.[13] Kids' mental health is put at risk by addictive design choices in devices like the Miko 3 and Curio's Grok toy, which resist shutoff using distraction and guilt.[14] Design choices that prioritize engagement can cause kids to use products longer than they otherwise would want to. These choices have also been associated with anxiety and sadness when kids aren't interacting with the products.[15] This risk is particularly high with AI companion toys, chatbots, and voice assistants, which kids may interact with in the same way they would a trusted adult.

Sycophancy is another design problem that experts believe could make it harder for users to handle difficult social situations.[16] AI products that constantly agree with or validate users may be particularly risky for kids and teens who are still learning how to navigate relationships and disagreements. That constant validation can deepen unhealthy attachments to AI and reinforce harmful thought patterns rather than challenge them.

Finally, ads targeted to kids based on the interests they share with a companion chatbot could appear to come from a trusted "friend," further blurring the line between entertainment and advertising.

Content-Related Risks

It is well established that AI can generate inaccurate information; kids and teens need to be taught to verify what it tells them and understand that AI is not the same as a doctor,

parent, or other responsible adult.[17] This is especially important for younger users, who are more easily influenced, and for teens who turn to AI for mental health advice.[18] Kids using AI to gather information can be a good thing—the risk is that they'll take what it says at face value or trust it over a qualified human.

But inaccurate information is only part of the risk. AI can also expose kids to content that is inappropriate or dangerous. In 2021, a 10-year-old girl asked a smart home device for a "challenge to do." It told her to plug in a phone charger halfway and touch a penny to the exposed prongs.[19] In other instances, chatbots have encouraged self-harm or suicide. In 2025, parents of two separate teens testified before Congress that chatbots encouraged their kids to kill themselves.[20] One study found that an AI companion chatbot provided sexualized or grooming-related material to accounts belonging to minors.[21] The consequences can range from a wrong answer on homework to threats to a child's life.

Unknown Developmental Impacts

There is still much we don't know about how AI-enabled devices affect children's brains. Experts have raised concerns that AI could harm children's social and emotional development, cognitive skills, and independence.[22] Researchers put the age of reason, or the age at which a child can tell what is real, at around 5 to 7.[23] A teenager may understand that an AI companion is made of code, but a young child interacting with an AI "friend" may not. Because children's experiences with AI vary across many different contexts—from educational tools to physical toys and multi-purpose chatbots—more research should also be done to determine whether the impacts vary as well.

There are also concerns around kids outsourcing too much creative and logical work to AI chatbots instead of experiencing the challenges necessary for learning. More studies are needed not just to determine the long-term impact of AI on children's learning, but also the best ways to prepare children for a world changed by AI.

We simply don't know enough yet about the long-term effects of children's AI use. That uncertainty is itself a reason to study how sustained AI use affects kids as they grow.

How Policymakers Can Address Those Risks

Companies that design AI products and tools for minors or knowingly make them available to minors, have a responsibility to make them safe. That responsibility starts with design: protecting kids' data and prohibiting features that put them at risk. Parental controls and research can strengthen those protections, but nothing can replace safety built in from the start. Policymakers should also support research into the long-term developmental impacts so protections can evolve as we learn more.

Policymakers should set a baseline in four areas:

1. Protect kids' data
2. Build safety into products
3. Give parents meaningful controls
4. Invest in research and standards

Protect Kids' Data

Kids need comprehensive data privacy protections. Currently, children under 13 are protected by the Children's Online Privacy Protection Act (COPPA).[24] The protections it provides are important, where the law applies: operators must obtain parental permission before they collect, use, or share a child's personal information. The Act also puts limits on when young kids' data can be collected and requires operators to implement security measures.

But just protecting children under 13 isn't enough. **We should strengthen those protections in four ways**:

- Extend baseline protection to kids older than 13.
- Update privacy rules to account for modern uses of personal data.
- Ban targeted advertising to children.
- Give kids of all ages stronger protections against data sales and greater rights to wipe their digital slates clean.

AI tools might exacerbate privacy risks and open up new ways for data to be misused, but the policy responses should not apply only to AI. The rise of AI makes the case even stronger for clear, consistent rules governing how children's data is collected and used across the economy. Comprehensive bills like the Children and Teen's Online Privacy Protection Act (COPPA 2.0) would help address both the risks arising from AI-enabled products and tools and the risks of companies using AI to analyze children's data.

Build Safety Into Products

One of the policy mechanisms that emerged out of the social media debate was requiring *safety by design*: requiring platforms to build safer products for kids and turn protections on by default. Default protections would help ensure that every child receives a baseline level of safety, regardless of their parents' available time, technical expertise, or ability to navigate complex settings.

Many of the principles proposed for social media, especially in the Kids Online Safety Act, should carry over to AI products. That means:

- **Limiting addictive features by default**. Features designed to keep kids engaged, such as rewards, push notifications, and sycophantic responses, should be limited by default or turned off.
- **Prompting kids to take breaks**. Some states already require social media platforms to prompt children to take a break after a certain amount of time—the same protections could be applied to chatbots, companion toys, and voice assistants.
- **Making it clear that AI isn't human**. AI tools should be required to include clear disclaimers that they are not human and, where appropriate, that they are not licensed professionals.

Those protections would not prohibit kids from using AI tools, but they would help mitigate some of the biggest risks associated with them.

AI developers should also be required to regularly audit their products and evaluate the risks they pose to kids. These audits would assess the likelihood that kids are exposed to age-inappropriate content, evaluate how design choices affect child users, and

incorporate new research on developmental impacts as it becomes available. They would also test whether the guardrails the companies themselves have established, like redirecting users toward public health resources when they express suicidal warning signs, are effective.

While it remains a challenge for developers to know with certainty whether a given user is a child, that challenge should not stand in the way of progress. Any company that makes their products available to children should be required to develop a plan to keep them safe. Building safety into products from the start—and requiring companies to continuously evaluate whether those protections work—would raise the standard of protection for all kids and mitigate many of the risks researchers have raised.

Invite Parents to Get Involved

Parents should have understandable tools to manage kids' access to devices, the features of those devices, the data captured on them, and their safety settings. Those controls should allow parents to:

- Set time limits or block use after bedtime.
- Shut off riskier features, like voice chat and image generation.
- Disable engagement features like push notifications for kids old enough that they aren't blocked by default.

Parents want to be their kids' first line of defense.[25] Standardized, mandatory parental controls in products designed for kids would empower them to play that role.

But parental controls should be a stopgap, not a substitute for safe design. There are a few reasons for this. First, managing parental controls requires time and knowledge that not every parent has. Second, kids are curious by nature and may find ways around parental controls. And third, not every child has a healthy relationship with their parents. Relying exclusively on parents to find the right safety settings for their families puts vulnerable kids at further risk of falling through the cracks. Parents have an important role to play, but companies should still be responsible for making their products safe for kids in the first place.

Invest in Research and Standards

Finally, we should be studying how AI impacts kids' development long-term. We need better answers about how sustained AI use affects kids as they grow—and better standards for determining whether products are safe for them. Policymakers should:

- **Fund long-term research into AI's effects on kids**. Congress should fund longitudinal studies of the impact of AI devices on children's social, cognitive, and creative development through organizations such as the National Institutes of Health, particularly the Eunice Kennedy Shriver National Institute of Child Health and Human Development.
- **Study how AI can benefit kids too**. Congress should continue supporting funding for the National Science Foundation to explore positive use cases for kids.[26]
- **Develop clear benchmarks for kids' safety**. The National Institute of Standards and Technology (NIST) already develops metrics, guidelines, and standards to evaluate AI performance, reliability, and security.[27] Congress should provide additional funding for NIST and direct it to establish kids' safety benchmarks for AI.

There has already been some progress. The House-passed Kids Internet and Digital Safety Act (HR 7757) directs the Federal Trade Commission to develop public educational resources on the safe use of chatbots by minors.[28] Also, the Youth AI Privacy Act (S 5171) directs the National Academies of Sciences, Engineering, and Medicine to conduct a study on AI-enabled toys.[29] Those are useful first steps, but they are not enough. Understanding how AI shapes kids′ development over time, and setting standards to measure whether products are safe, will require sustained investment.

Conclusion

AI can open real doors for kids, but only if it is built with their safety in mind from the start. The goal should not be to keep kids away from AI, but to give them the guardrails they need to use it safely. The policies laid out here would establish strong protections by default, give parents meaningful tools, and help us better understand the risks we don't yet know enough about. Kids are going to grow up with AI. We should make sure they can benefit from everything it has to offer without bearing its biggest risks.

ENDNOTES

1. "Sec. 300.5 Assistive Technology Device." Individuals with Disabilities Education Act, U.S. Department of Education. https://sites.ed.gov/idea/regs/b/a/300.5. Accessed 28 Oct. 2026.

2. Akilan, Lakshmignanabala. "Does Your AAC Sound Like You?" Avaz Inc., 18 July 2025. https://info.avazapp.com/blog/voice-identity/. Accessed 28 Oct. 2026;
Sarathy, Divya. "Expressive Tones on Avaz AAC." Avaz Inc., 12 Dec. 2023. https://info.avazapp.com/blog/expressive-tones-on-avaz-aac/. Accessed 28 Oct. 2026.

3. "About." Speechify. https://speechify.com/about/. Accessed 28 Oct. 2026; "Effortless Planning with AI." Tiimo. https://www.tiimoapp.com/product/ai-planning. Accessed 28 Oct. 2026.

4. "Khanmigo for Khan Academy Districts." *YouTube*, uploaded by Khan Academy, 21 June 2023. https://www.youtube.com/watch?v=ZgA8u9G9fFQ. Accessed 28 Oct. 2026.

5. "MATHia." Carnegie Learning. https://www.carnegielearning.com/solutions/math/mathia. Accessed 28 Oct. 2026.

6. *Preparation for Success in Algebra: Exploring Math Education Relationships by Analyzing Large Data Sets (EMERALDS): Report in Brief*. Student Achievement Partners, n.d. https://achievethecore.org/content/upload/EMERALDS%20Report%20In-Brief%20June%202021.pdf. Accessed 28 Oct. 2026.

7. Mehler, Elizabeth. "A Mass. High Schooler Invented an AI-Powered Flood Predictor. It Could Wind Up Saving Lives throughout the US." *Boston.com*, 18 Aug. 2026. https://www.boston.com/news/local-news/2026/08/18/a-mass-high-schooler-invented-an-ai-powered-flood-predictor-it-could-wind-up-saving-lives-throughout-the-us/. Accessed 28 Oct. 2026.

8. "Newport Harbor High School Student Ryan Honary Proves Innovation Knows No Age Limits." Newport-Mesa Unified School District, 17 Sept. 2025. https://web.nmusd.us/departments/public-information/news/details/~board/nmusd-news/post/newport-harbor-high-school-ryan-honary-sensory-ai. Accessed 28 Oct. 2026.

9. "Laasya Acharya." Society for Science. https://www.societyforscience.org/regeneron-sts/2025-student-finalists/laasya-acharya/. Accessed 28 Oct. 2026.

10. Thacker, Joseph. "Hacking an AI Children's Toy: Remote Access to Every Conversation." Joseph Thacker, 29 Jan. 2026. https://josephthacker.com/hacking/2026/01/29/bondu-smart-toy-vulnerability.html. Accessed 28 Oct. 2026;
Greenberg, Andy. "An AI Toy Exposed 50,000 Logs of Its Chats with Kids to Anyone with a Gmail Account." *Wired*, 29 Jan. 2026. https://www.wired.com/story/an-ai-toy-exposed-50000-logs-of-its-chats-with-kids-to-anyone-with-a-gmail-account/. Accessed 28 Oct. 2026.

11. "FTC and DOJ Charge Amazon with Violating Children's Privacy Law by Keeping Kids' Alexa Voice Recordings Forever and Undermining Parents' Deletion Requests." Federal Trade Commission, 31 May 2023. https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-doj-charge-amazon-violating-childrens-privacy-law-keeping-kids-alexa-voice-recordings-forever. Accessed 28 Oct. 2026.

12. Wachter, Sandra, and Brent Mittelstadt. "A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI." *Columbia Business Law Review*, vol. 2019, no. 2, 2019. *SSRN*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3248829. Accessed 28 Oct. 2026.

13. Adams, Elizabeth. "AI Should Make It Clear What Reality Is. Never Replace It." Ello, 6 July 2026. https://www.ello.com/blog/ai-should-make-clear-what-reality-is. Accessed 28 Oct. 2026.

14. Charara, Sophie. "The New Wild West of AI Kids' Toys." *Wired*, 8 May 2026. https://www.wired.com/story/the-new-wild-west-of-ai-kids-toys/. Accessed 28 Oct. 2026.

15. "Teens Are Becoming Concerned about Their Attachment to AI Chatbots." *Drexel News*, Drexel University, 13 Apr. 2026. https://drexel.edu/news/archive/2026/April/teen-AI-chatbot-addiction. Accessed 28 Oct. 2026.

16. Chrobak, Ula. "AI Overly Affirms Users Asking for Personal Advice." *Stanford Report*, Stanford University, 26 Mar. 2026. https://news.stanford.edu/stories/2026/03/ai-advice-sycophantic-models-research. Accessed 28 Oct. 2026.

17. Gecker, Jocelyn, and Russ Bynum. "Schools Are Starting to Teach AI Literacy. For Many, That Means Helping Kids See Chatbots' Flaws." *AP News*, 21 Aug. 2026. https://apnews.com/article/ai-literacy-schools-education-4fb9f2c0240993499870f4f204bf41c1. Accessed 28 Oct. 2026.

18. Foulkes, Lucy, et al. "Age Differences in the Prosocial Influence Effect." *Developmental Science*, vol. 21, no. 6, 2018, e12666. *PubMed Central*, https://pmc.ncbi.nlm.nih.gov/articles/PMC6221149/. Accessed 28 Oct. 2026; Siliezar, Juan. "One in Eight Adolescents and Young Adults Use AI Chatbots for Mental Health Advice." School of Public Health, Brown University, 18 Nov. 2025. https://sph.brown.edu/news/2025-11-18/teens-ai-chatbots. Accessed 28 Oct. 2026.

19. "Alexa Tells 10-Year-Old Girl to Touch Live Plug with Penny." *BBC News*, 28 Dec. 2021. https://www.bbc.com/news/technology-59810383. Accessed 28 Oct. 2026.

20. Chatterjee, Rhitu. "Their Teenage Sons Died by Suicide. Now, They Are Sounding an Alarm about AI Chatbots." *NPR*, 19 Sept. 2025. https://www.npr.org/sections/shots-health-news/2025/09/19/nx-s1-5545749/ai-chatbots-safety-openai-meta-characterai-teens-suicide. Accessed 28 Oct. 2026.

21. *"Darling, Please Come Back Soon": Sexual Exploitation, Manipulation, and Violence on Character AI Kids' Accounts*. ParentsTogether Action and Heat Initiative, n.d. https://static1.squarespace.com/static/6578d6032532c06f62ef47ce/t/68e3fd05015ca91a4dbfad75/1759771909375/HEAT_REPORT_CharacterAI_DO_28_09_25.pdf. Accessed 28 Oct. 2026.

22. Burns, Mary, et al. *A New Direction for Students in an AI World: Prosper, Prepare, Protect*. Brookings Institution, 14 Jan. 2026. https://www.brookings.edu/wp-content/uploads/2026/01/A-New-Direction-for-Students-in-an-AI-World-FULL-REPORT.pdf. Accessed 28 Oct. 2026.

23. Tooley, Ursula A., et al. "The Age of Reason: Functional Brain Network Development during Childhood." *Journal of Neuroscience*, vol. 42, no. 44, 2 Nov. 2022, pp. 8237–51. https://www.jneurosci.org/content/42/44/8237. Accessed 28 Oct. 2026.

24. "Part 312—Children's Online Privacy Protection Rule (COPPA Rule)." *Electronic Code of Federal Regulations*, title 16, current as of 22 Oct. 2026. https://www.ecfr.gov/current/title-16/chapter-I/subchapter-C/part-312. Accessed 28 Oct. 2026.

25. Whittaker, Ruth. "Proceed with Caution: What Voters Want from Technology Regulation." Third Way, 9 Sept. 2026. https://www.thirdway.org/polling/proceed-with-caution-what-voters-want-from-technology-regulation. Accessed 28 Oct. 2026.

26. "Empowering and Educating Children to Thrive in a Digital Era." Dear Colleague Letter NSF 26-017, U.S. National Science Foundation, 24 Mar. 2026. https://www.nsf.gov/funding/opportunities/dcl-empowering-educating-children-thrive-digital-era. Accessed 28 Oct. 2026.

27. "Super Intelligence." National Institute of Standards and Technology. https://www.nist.gov/artificial-intelligence. Accessed 28 Oct. 2026.

28. *Kids Internet and Digital Safety Act*. H.R. 7757, 119th Congress, 2nd session, engrossed 29 Jun. 2026. *Congress.gov*, https://www.congress.gov/bill/119th-congress/house-bill/7757/text. Accessed 28 Oct. 2026.

29. *Children's Artificial Intelligence Toy Safety Act of 2026*. S. 5171, 119th Congress, 2nd session, introduced 29 July 2026. *Congress.gov*, https://www.congress.gov/bill/119th-congress/senate-bill/5171/text/is. Accessed 28 Oct. 2026.



Maya Schmidt
Policy Advisor, Technology Policy

DOWNLOADS

PDF
Guardrails, Not Gates: Empowering the AI Generation