

PUBLIC SUBMISSION

Received: May 29, 2025 Tracking No. mba-2vk3-kqxx Comments Due: May 28, 2025 Submission Type: API
--

Docket: NSF-2025-OGC-0001
NITRD_FRDOC_0001

Comment On: NSF-2025-OGC-0001-0001
Request for Information: Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

Document: NSF-2025-OGC-0001-DRAFT-0308
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Children and Screens: Institute of Digital Media and Child Development

General Comment

See attached file(s)

Attachments

Docket ID No. NSF-2025-OGC-0001 Comments of Children and Screens_ Institute of Digital Media and Child Development

Docket ID No. NSF-2025-OGC-0001 Comments of Children and Screens - Institute of Digital Media and Child Development

Docket ID No. NSF-2025-OGC-0001 Comments of Children and Screens- Institute of Digital Media and Child Development

**Comments of Children and Screens: Institute of Digital Media and Child Development
to the Networking and Information Technology Research and Development National
Coordination Office Docket ID No. NSF-2025-OGC-0001**

Children and Screens: Institute of Digital Media and Child Development (“Children and Screens”) appreciates the opportunity to submit comments to the Networking and Information Technology Research and Development (NITRD) National Coordination Office’s (NCO) request for information on artificial intelligence (AI). The request’s focus—research and development (R&D) to accelerate AI-driven innovation, enhance U.S. economic and national security, promote human flourishing, and maintain the United States’ preeminence in AI—is admirable. Since 2013, Children and Screens has worked to help young people lead healthy lives in a digital world by synthesizing and disseminating the latest scientific research, while also supporting advancements in the field through research funding and scientific convenings.

AI, particularly generative AI like ChatGPT and DeepSeek¹, is poised to make profound consequences for multiple industries, education, and social life. Children will necessarily be impacted by such disruptive technologies. The Trump-Vance Administration recognizes the opportunities AI advances offer, illustrated by executive orders creating pathways for AI innovation and integration into educational systems². AI indeed presents monumental potential, but does so along with significant risk. As with other key sectors, such as energy and biosciences, harms must be mitigated through rigorous research to maximize the numerous benefits of AI and avoid pitfalls.

The rapid development of powerful AI models is outpacing policy and programmatic efforts to ensure these systems are aligned with human values and safe - especially for children. Current alignment methods³ are technically inadequate, poorly resourced, and unlikely to evolve alongside AI capabilities. Without significant public investment and regulation, this imbalance may result in societal and individual harms and lead to more stifling regulation down the road. Children, who will live longest with the consequences of today’s choices, are especially at risk. Ensuring their welfare must be central to alignment and safety research, which demands Executive action and long-term public sector commitment.

Research investments will maintain U.S. advantage and global competitiveness

¹ Chatgpt, <https://chatgpt.com/> (last visited May 28, 2025).; 深度求索 DeepSeek, <https://www.deepseek.com/> (last visited May 28, 2025)

² Exec. Order No. 14179, 90 FR 8741 (2025); Exec. Order No. 14277 90 FR 17519 (2025)

³ *Alignment* refers to the process of ensuring AI systems act in ways consistent with human intentions, values, and societal norms.; Tom Duenas & Diana Ruiz, *The Ethics of Character Engineering in Artificial Superintelligence* (unpublished manuscript, 2024)

The United States must lead in AI innovation, but innovation alone is not enough. Global leadership and maintaining the American technical advantage will depend on the trustworthiness and safety of AI systems. If the U.S. rushes AI products to deployment in the military, educational, or commercial sectors before adequate safety and alignment can be assured, it risks ceding competitive advantage to other actors.

Policy experts have emphasized the importance of responsible deployment to avoid innovation-stifling regulation in this sector⁴. The key to avoiding heavy-handed regulation is to strategically invest in safety and alignment research at the outset of this technological development. The U.S. must therefore proactively invest in fundamental R&D around AI. Given that today's children will inherit the consequences of current AI advances, this R&D must be designed with their safety and well-being at its core. Centering young people in these efforts will not only protect their development and learning in the short term, but also ensure they are equipped to thrive in — and lead — the AI-driven world of tomorrow. We recommend investments in three critical areas: alignment and safety, AI developed for children, and k-12 education.

The need for foundational alignment and safety research

Rapidly evolving AI frontier model development and AI system deployment nearly ensures serious alignment and safety issues will emerge at some point. Frontier model misalignment and safety failures from future systems pose serious risks⁵ not only to children, but to public trust and the technology's ability to have positive educational and other impacts as envisioned by the Administration. Addressing AI safety concerns is currently substantially less profitable than improving AI capacity, so frontier model developers primarily pursue AI capacity development over addressing safety, or laying a foundation for the technology's value for learning and more. This incentive structure is leading to a future in which models are immensely powerful and complicated, and deployed everywhere - especially around children - but with limited safety or utility.

A thoughtful, evergreen approach is needed to assure safe and prosperous AI development over the long run. Thought leaders in this space recognize the need for the public sector to play a role in securing basic, foundational alignment and AI safety research⁶ so that foundational AI systems are able to safely fill the pivotal role they will likely play in the nation's future. Major companies themselves note these safety insufficiencies and their executives are calling for regulation⁷.

⁴ James Pethokoukis, *Permissionless Innovation in the Age of AI*, Am. Enter. Inst. (Mar. 5, 2024), <https://www.aei.org/articles/permissionless-innovation-in-the-age-of-ai/>

⁵ Dan Hendrycks et al., *Aligning AI With Shared Human Values* (unpublished manuscript, Aug. 2021)

⁶ Aaron Gregg, Cristiano Lima-Strong & Gerrit De Vynck, *AI Poses Risk of 'Extinction,' Industry Leaders Warn*, Wash. Post (May 30, 2023), <https://www.washingtonpost.com/technology/2023/05/30/ai-extinction-risk/>

⁷ Univ. of Oxford, *World Leaders Still Need to Wake Up to AI Risks, Say Leading Experts Ahead of AI Safety Summit*, Univ. of Oxford (May 20, 2024), <https://www.ox.ac.uk/news/2024-05-20-world-leaders-still-need-wake-ai-risks-say-leading-experts-ahead-ai-safety-summit>; Cecilia Kang, *OpenAI's Sam Altman Urges A.I. Regulation in Senate Hearing*, N.Y. Times (May 16, 2023), <https://www.nytimes.com/2023/05/16/technology/openai-altman-artificial-intelligence-regulation.html>

Today's approaches to technical AI alignment have little prospect of remaining evergreen as AI capacity increases.

When critical safety research is not being addressed by markets, public sector research is essential, particularly to generate the evidence needed by Congress and other policy leaders.

Unless child welfare and flourishing becomes a core component of all alignment and safety research, AI poses both immediate and long-term risks to children. Congress has already acknowledged some of these near-term harms in passing the TAKE IT DOWN Act⁸, which will help victims of AI-generated child sexual abuse materials (CSAM). Recent incidents signal additional severe harms not addressed by current legislation. For example, a 14-year-old boy died by suicide after engaging with an AI character inspired by fiction and his own preferences on the character.ai platform⁹. Since AI systems do not develop their outputs in the same way as humans, alignment with societal values is not guaranteed. Broader safety testing must consider the needs of child users interacting with AI systems. Child users will interact with AI systems more frequently over time, and often with greater compulsion involved, as AI is integrated even more deeply into schools and entertainment platforms. To assure rigorous safety standards, children must be prioritized in fundamental safety research.

The need to design AI for children

Another fundamental gap in the current AI landscape involves AI made specifically for children. The history of digital media has shown that children are uniquely vulnerable to the risks and dangers associated with technological developments¹⁰ because they lack the developmental capacity, knowledge, and experience required to use them appropriately. Generative AI is a clear concern in this regard, but all forms of AI must also be included. This is especially true within domains like social media, which deploy machine learning techniques to engage child users in a manner that conflicts with children's optimal development. Newer AI systems are unlikely to be different in this regard. Risks to children will increase proportional to AI capacity development and commercial and educational deployment.

AI is already being deployed around children at a rate disproportional to the AI deployment rates targeting adults. Dozens of AI products marketed as “child-friendly” are commercially available,

⁸ S. 146, 119th Cong. (2025)

⁹ Kate Payne, *An AI Chatbot Pushed a Teen to Kill Himself, a Lawsuit Alleges*, AP News (Oct. 25, 2024), <https://apnews.com/article/9d48adc572100822fdb3c90d1456bd0>

¹⁰ Jennifer J. Chen & Jasmine C. Lin, *Artificial Intelligence as a Double-Edged Sword: Wielding the POWER Principles to Maximize Its Positive Effects and Minimize Its Negative Effects*, 25 *Contemp. Issues Early Child.* 146 (2023); Sonia Livingstone & Leslie Haddon eds., *Kids Online: Opportunities and Risks for Children* (Policy Press 2009); Ellen A. Wartella & Nancy Jennings, *Children and Computers: New Technology, Old Concerns*, *Future Child.* (2000); Ian Hutchby & Jo Moran-Ellis, *Children, Technology and Culture: The Impacts of Technologies in Children's Everyday Lives* (Routledge 2013); Ann-Marie Kennedy, Katherine Jones & Janine Williams, *Children as Vulnerable Consumers in Online Environments*, 53 *J. Consumer Aff.* 1478 (2019)

and for free in many cases¹¹. For children in schools, participation with AI programs may be more-or-less compulsory, such as Google Gemini being integrated into Google classroom and Google workspace¹².

AI is becoming an increasingly prevalent environmental factors joining the host of environmental and biological factors that interact across development to shape how children learn, act, socialize, and grow¹³. Over a century of developmental science has shown that children are not simply “little adults”. Children have distinct developmental needs of their own, which change as they progress through developmental stages. Even simple developmental differences - like the fact that most adults work at jobs, whereas most children attend school - make the implicit demands of users’ lives vastly different. For example, working adults might be better served by an AI that directly answers questions in completely factual ways, whereas a child might be better served by an AI that guides them through the thinking process toward an answer.

This disconnect between users' needs raises significant AI suitability concerns. A Large Language Model (LLM) trained on adult language data and fine-tuned on adult human feedback should not be framed as suitable for widespread deployment around children. AI systems might inadvertently disrupt our traditional learning processes, such as cognitive offloading, but also in ways that are currently not known, due to insufficient research. These learning processes could have long-term consequences on the entire education system, as well as prospects for future employability and workforce training. Interacting with children in a developmentally appropriate way is a non-trivial technical problem to solve. Research is desperately needed in this area, especially given that the topic receives little focus from major AI companies.

Technically speaking, current AI systems designed for children are third party “AI products,” rather than child-oriented systems designed by frontier model developers. This distinction includes products with wide adoption¹⁴. AI products, while sometimes innovative and useful, tap into frontier models trained on adult content reflecting adult world-views.

The approach taken by “AI products” - retrofitting frontier models with ad-hoc “child filters” - is therefore not a viable long-term solution to child safety. These methods are generally driven by superficial content moderation heuristics, and lack the ability to reconfigure and reorganize knowledge represented in the underlying model in a manner more suitable for a child user.

¹¹ For e.g., see *Amazon Echo dot for Kids*, Amazon's smart speaker specifically designed for children; *Miko*, an AI companion robot for kids; *Luka*, a reading robot companion for Kids; *Cozmo*, an educational toy robot; *CharacterAI*, a chatbots children can use

¹² For a deeper discussion on these issues see Mathilde Neugnot-Cerioli & Olga Mu. Laurenty, *The Future of Child Development in the AI Era: Cross-Disciplinary Perspectives Between AI and Child Development Experts* (May 2024) (unpublished manuscript)

¹³ *Id.* (should be Neugnot-Cerioli if moved)

¹⁴ For e.g. see *KhanMigo*, Khan Academy's AI-powered tutor; *Khan Academy*.

<https://blog.khanacademy.org/announcing-free-khanmigo-for-select-arizona-districts/>

Children are already being impacted in negative ways, and in some cases irreparable harm has been done, as in the aforementioned case of suicide¹⁵. Practically speaking, using research to design AI systems from the ground up with children in mind is the best way to meet their needs. Children are more vulnerable than adults to various types of misinformation. Current AI systems are trained on data reflecting adult norms, values, and knowledge, which might distort a child user's sense of self, values formation, and relationship formation, especially with repeated exposure¹⁶.

As AI systems grow more sophisticated and autonomous, the gap between their capabilities and children's vulnerabilities is likely to widen—making child-centered AI design a critical priority. To ensure AI is safe for children in the long term, developers must be compelled or incentivized to prioritize children's needs from the outset. This requires more robust research into how children interact with and understand current AI systems, and in turn are impacted by these systems.

Safe and effective AI in k-12 education

Education is critical to children's safe and effective use of AI¹⁷. AI can facilitate learning, but it must be applied appropriately¹⁸. As a disruptive technology, it is imperative that children learn to use AI safely and proficiently. Professional development for educators is key to this goal¹⁹.

By exposing children to AI concepts and AI itself, educators and caregivers can build children's competencies to prepare them to be competitive in an AI-fueled economy and workforce. However, deploying AI and AI education in the classroom, without care, will only hurt our youth by disrupting their development and hindering their learning in important domains. This puts them at a disadvantage as they mature into adulthood and it disturbs their ability to compete in the future²⁰.

¹⁵ Payne *supra* note 9; Nomisha Kurian, *AI's Empathy Gap: The Risks of Conversational Artificial Intelligence for Young Children's Well-Being and Key Ethical Considerations for Early Childhood Education and Care*, 26 *Contemp. Issues Early Child.* 132 (2023); Eduard Fosch-Villaronga et al., *Toy Story or Children Story? Putting Children and Their Rights at the Forefront of the Artificial Intelligence Revolution*, 38 *AI & Soc'y* 133 (2023); Judith H. Danovitch et al., *Children's Understanding and Use of Voice-Assistants: Opportunities and Challenges*, in *Handbook of Children and Screens* 1 (Dimitri A. Christakis & Lauren Hale eds., Springer 2025)

¹⁶ Katie Notopoulos, *They Bought an \$800 AI Robot for Their Kids. Now the Company Is Shutting Down — and Children Are Having to Say Goodbye*, *Bus. Insider* (Dec. 11, 2024), <https://www.businessinsider.com/moxie-robot-toy-shutting-down-kids-embodied-goodbye-2024-12>; Radhika Garg & Subhasree Sengupta, *He Is Just Like Me: A Study of the Long-Term Use of Smart Speakers by Parents and Children*, 4 *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, Art. 11 (2020)

¹⁷ Exec. Order No. 14277 90 FR 17519 (2025)

¹⁸ Ying Xu et al., *Growing Up with Artificial Intelligence: Implications for Child Development*, in *Handbook of Children and Screens* 611 (Dimitri A. Christakis & Lauren Hale eds., Springer 2025)

¹⁹ Exec. Order No. 14277 90 FR 17519 (2025); Zu et al., *supra* note 18 Ying Zu chapter of handbook

²⁰ Neugnot-Ceroli & Laurenty *supra* note 12

Being transparent with children about the nature of AI, and helping them establish and maintain boundaries with this technology, is essential for productive use of AI²¹. Research has found that children interact with AI fundamentally differently than they do with actual adults—this underlies why human adults remain better instructors than AI²². Yet understanding this difference also means we can design and deploy AI in education to complement learning from human educators, and even enhance said learning²³. As noted in the preceding section, children may be better served by an AI that guides them through the thinking process, rather than provide a factually correct answer. When AI is used to support specific learning outcomes, this distinction is critical. For example, Khan Academy Kids, which uses a responsive-learning algorithm to guide them through their work²⁴. In contrast, common LLMs, like ChatGPT, can simply perform assignments for students, providing no real learning benefits.

In any education setting, the use of AI-based prediction and decision making should be approached with extreme caution. AI prediction and decision making is not inherently superior to human prediction and decision making, and in fact can be far inferior in certain contexts²⁵. AI decision making models are also unable to accurately consider important life changes, or contextual factors in learning, and development. Using models as the sole predictor of student outcomes will miss some of the more subtle signs of progress evident in the classroom²⁶. AI decision making systems can also encourage individuals to align their performance with arbitrary metrics of the system, rather than encourage broad knowledge acquisition and skill building²⁷. Finally, AI decision making systems without human oversight offer no recourse when errors are made, as AI systems are predictive “black box” systems that produce outputs based on knowledge representation systems. Consequently, their decision making processes can be uninterpretable in ways that human decision-making is not. To safeguard against these risks, any predictive or decision-making AI system deployed in educational settings must first be adequately tested safely in real-world settings. AI systems should never be deployed under the assumption that they are superior or more “objective” to human actors, and should only be deployed in the least consequential circumstances. Humans must be involved in and oversee all AI-driven decision making, and a human must be responsible for any decisions made. For

²¹ Ying Xu, *Different but Complementary: Navigating AI’s Role in Children’s Learning and Development*, Joan Ganz Cooney Ctr. (Oct. 7, 2024), <https://joanganzcooneycenter.org/2024/10/07/different-but-complementary/>

²² *Id.*

²³ For e.g. *When tutoring, Intelligent Tutoring Systems can be as effective at scaffolding and providing step-based tutoring*; Kurt VanLehn, *The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems*, 46 Educ. Psychol. 197 (2011)

²⁴ Khan Academy, *Khan Academy Kids*, <https://learn.khanacademy.org/khan-academy-kids/> (last visited May 29, 2025)

²⁵ Angelina Wang et al., *Against Predictive Optimization: On the Legitimacy of Decision-Making Algorithms That Optimize Predictive Accuracy*, 1 ACM J. Responsib. Comput. Art. 9 (2024); Alexandre Chiavegatto Filho et al., *Data Leakage in Health Outcomes Prediction With Machine Learning*, 23 J. Med. Internet Res. e10969 (2021)

²⁶ Wang et al., *supra* note 25; Filho et al., *supra* note 25; BBC News, *A-Levels and GCSEs: How Did the Exam Algorithm Work?*, BBC (Aug. 20, 2020), <https://www.bbc.com/news/explainers-53807730>

²⁷ Alene Rhea et al., *Resume Format, LinkedIn URLs and Other Unexpected Influences on AI Personality Prediction in Hiring: Results of an Audit*, in *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* 572 (ACM 2022)

instance, an AI system could be used to identify academic strengths and weaknesses of students based on test scores, but any results should be vetted by a trained educator or counselor. The results should be considered just one tool among many when the educator or counselor is deciding how to best help a student. The results alone should *not* be used to inform high-consequence decisions, like academic placement or career advice. An appropriate use could include informing personalized curricula for students. Regardless of the use-case, a human should be held responsible for the ultimate decision, and the insights of a trained professional should always be trusted over those produced by an AI system.

The EDSAFE AI Alliance provides 4 guiding principles to help direct the integration of AI into k-12 education²⁸. The principles can apply to any type of AI being considered, including generative AI and AI decision-making systems. To be safe and effective, AI use and education should be:

- *Safe* - this includes every aspect of safety, from social development to data privacy
- *Accountable* - establish benchmarks that are collaboratively defined by a diverse group of constituents and require transparency
- *Fair* - AI products and education should be accessible by all students, and free from bias
- *Effective* - Carefully consider the utility of AI in education and establish methods to assess efficacy

Recommendations and Conclusions

Children must not be an afterthought in our national AI research strategy. Embedding child-centered design principles into foundational AI development—from training data to deployment contexts—will ensure these systems contribute to human flourishing rather than undermine it. It is essential that the U.S. lead not only in AI innovation and competitiveness, but in responsible stewardship that reflects our nation’s values and commitment to future generations. To achieve this, we propose the following recommendations:

The Need for Foundational Alignment and Safety Research

- **Fund research into evergreen AI alignment methods that include child welfare as a core consideration.**
 - Research child-inclusive AI safety benchmarks and risk assessment frameworks.
 - Establish empirical groundwork for pre-deployment safety testing of AI systems likely to be used by or around children.
 - Ensure that AI R&D addresses risks mitigation specific to children.
 - Encourage collaboration between AI researchers and child development experts.
 - Study and develop ethical use and data protection standards for child-facing AI.
- **Advance legislation that supports research into AI-driven harms.**

The Need to Design AI for Children

- **Ensure innovation aligns with child health and well-being.**

²⁸ EDSAFE AI Alliance, *SAFE Benchmarks Framework*, EDSAFE AI, <https://www.edsafeai.org/safe> (last visited May 28, 2025)

- Support the research and development of dedicated child-focused AI systems.
- Compel or incentivize frontier model developers to integrate child-centered design from the earliest stages:
 - Require that AI training data reflect varied and developmentally appropriate content for children.
 - Supervised labeling of training data.
 - Reinforcement learning protocols that consider children and their development.
- **Promote transparency and AI literacy with families and children about when and how AI is used in child-facing applications²⁹.**
- **Establish ethical guidelines for AI interactions with children based on developmental and educational psychological science.**
 - Invest in the research needed to inform these guidelines, including research on developmental outcomes from repeated AI interaction

AI in K–12 Education

- **Fund professional development for educators on AI literacy, integration, and student support.**
- **Critically evaluate areas of opportunity and risk before deployment of AI in different aspects of education.**
 - Invest in research on how children interact with AI in classroom settings, its impacts on learning and what might promote healthy, productive engagement for lifelong learning
 - Introduce AI education in developmentally appropriate ways that encourage learning, not dependency.
 - Support AI tools that complement—not replace—human teaching relationships. For example, tutor chatbots and intelligent tutoring systems that complement in-class learning
- **Adopt the EDSAFE AI Alliance's 4 guiding principles for school-based AI³⁰:**
 - Safe – Prioritize emotional, social, and data privacy protections.
 - Accountable – Ensure transparency and include diverse stakeholders in oversight.
 - Fair – Promote equitable access and proactively mitigate bias.
 - Effective – Evaluate AI's educational utility and impact on learning outcomes.
- **Exercise extreme caution when integrating AI prediction and decision making systems:**
 - All systems must be tested across varied populations in real-world environments before being deployed for any purpose
 - AI systems should not be deployed under the assumption that they are superior to human actors
 - Systems must be deployed to inform the least consequential decisions possible
 - There must be human oversight throughout the decision making process
 - A human must be responsible for any and all decisions informed by an AI system

²⁹ For e.g. YouTube algorithms

³⁰ EDSAFE AI Alliance *supra* note 28

We urge the NITRD and its partners to adopt these recommendations and center children in Federal AI R&D planning by investing in alignment and safety research, supporting educator training, and promoting child-specific AI design standards. With bold leadership and thoughtful policy, we can harness AI's potential while safeguarding the well-being of our most vulnerable and valuable population.

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the 2025 National AI R&D Strategic Plan and associated documents without attribution.

Prepared and submitted, this the 29th day of May 2025, by:

John Mootz, PhD
Senior Research and Policy Manager

Kris Perry, MSW
Executive Director

Daniel Hipp, PhD
Senior Research Coordinator

Reg Leichty, Esq.
Policy Advisor

Nandini Asavari Bharadwaj, MA
Impact Fellow

Children and Screens: Institute of Digital
Media and Child Development
276 5th Avenue
Suite 704, Office #900
New York, NY 10001

Colleen Kraft, MD, MBA, FAAP
National Scientific Advisory Board

Dimitri Christakis, MD, MPH
Chief Science Officer