Docket ID: ED-2026-IES-3268

Public Comment on Priority Research Questions for Evaluating IDEA Implementation

General-Purpose AI as Assistive Technology for Children with Disabilities

Submitted by Basil C. Puglisi, MPA, independent practitioner, basilpuglisi.com

To: Institute of Education Sciences, National Center for Education Evaluation and Regional Assistance, in response to the Request for Information published November 13, 2026 (91 FR 73218)

Summary

General-purpose AI can function as assistive technology when its use reduces a disability-related functional barrier and increases, maintains, or improves a person's functional capability. IDEA already defines an assistive technology device by that function, and the definition reaches items "acquired commercially off the shelf" (20 U.S.C. § 1401(1)). This comment asks NCEE to evaluate whether general-purpose AI reaches children with disabilities as assistive technology under IDEA. It also asks whether school and state AI policies block that use where an IEP calls for it.

The comment answers Questions 1, 2, 3, 6, 7, and 8. It draws on three working papers by the author, listed below, and cites each source by a live address.

The Basis for This Comment

The first paper, When AI Becomes Assistive Technology, names a functional category the author calls Artificially Intelligent Disability Assistance (AIDA). It proposes a three-prong functional test for when a particular use of general-purpose AI qualifies. A separate two-prong screen, drawn from Checkpoint-Based Governance, asks how authority and accountability are allocated once a use qualifies, and whether the use delivers a functional benefit (Puglisi, 2026a).

The second paper, When AI Policy Becomes an Accessibility Barrier, classifies six types of AI policy risk and maps which statutes reach which actors. Its matrix places kindergarten through grade 12 under IDEA and Section 504, and it closes with a model accommodation clause any institution can adopt (Puglisi, 2026b).

The third paper, A Call for an AI Disability Standard, holds that the protections against punishing people for using AI as a disability tool already exist. It proposes a dated notice and a validation layer modeled on systems already in use, such as disability parking placards (Puglisi, 2026c).

Question 1: The Most Pressing Areas to Evaluate Now

The three papers bear most directly on three of the six areas: access, procedural safeguards, and fiscal and resource effectiveness. They also raise questions under appropriate services and educational outcomes and under statewide systems, accountability, and flexibility, which Question 2 takes up.

Access. The IDEA definition asks what an item does for the child, not who designed it or how it was marketed. The AIDA Functional Test asks the same question in three prongs (Puglisi, 2026a):

1. A disability creates an identifiable functional barrier.
2. A particular AI capability directly reduces or bypasses that barrier.
3. The intervention increases, maintains, or improves the person's functional capability.

The third prong borrows the statutory verbs directly. IDEA requires each IEP team to "consider whether the child needs assistive technology devices and services" (20 U.S.C. § 1414(d)(3)(B)(v)). The first evaluation question is whether that consideration now reaches general-purpose AI, and on what basis teams accept or reject it.

Procedural safeguards. School AI policies written to protect assessment can reach the same functions a child with a disability needs for access. The second paper names six ways a policy reaches an assistive user (Puglisi, 2026b). They are use prohibition, process restriction, disclosure requirement, distribution or opportunity penalty, discipline or exclusion, and perception enforcement. Only the first stops the use, while the other five regulate how AI use is carried out, disclosed, received, or punished. Perception enforcement can reach a student whether or not AI was used, because detection tools act on whether work seems AI-made. Research on GPT detectors found that they "consistently misclassify non-native English writing samples as AI-generated" and "may unintentionally penalize writers with constrained linguistic expressions" (Liang et al., 2023). The open question for IDEA is what that mechanism means for a child whose writing is formulaic because of a disability, or restructured by an assistive tool.

Fiscal and resource effectiveness. IDEA counts "purchasing, leasing, or otherwise providing for the acquisition of assistive technology devices" as an assistive technology service (20 U.S.C. § 1401(2)(B)). The first paper's review identified no single general rule for who pays for a general-purpose AI subscription used as assistive technology. It also names vendor dependency, since a model retirement, a policy change, or a change in price or capability can remove an accommodation overnight (Puglisi, 2026a).

Question 2: Research Questions to Prioritize Over the Next Five Years

The questions below adapt the open research agendas of the first two papers to the IDEA areas.

Access

1. When IEP teams consider assistive technology, how often do they consider general-purpose AI tools, and what reasons do they record when they accept or reject one?
2. How consistently can different evaluators apply a functional test, such as the AIDA Functional Test, to the same use by the same child?
3. How many district and state AI policies carry an accommodation exception, or a cross-reference to the IEP and Section 504 processes?

Appropriate services and educational outcomes

4. Which functional barriers show benefit from general-purpose AI, and which show harm?
5. Where does dependence emerge, and which tasks should stay with the child?

6. Do functional gains measured in the child persist across grades and transitions?

Procedural safeguards

7. How do schools treat false-positive AI detection involving students with disabilities, and does a human with the time, the evidence, and the authority to overturn the flag review it before any discipline?
8. Does a requirement to disclose AI use produce different consequences for students with disabilities than for their peers?

Statewide systems, accountability, and flexibility

9. When a state or local agency uses AI to triage evaluations, eligibility, or service requests, who reviews its output, and what record shows individualized review?

Fiscal and resource effectiveness

10. Who pays for a general-purpose AI subscription used as assistive technology under an IEP, and what happens to the child's services when the vendor changes its price, its policy, or its model?

Question 3: How These Questions Support IES's Statutory Responsibilities

The notice describes the IES role under IDEA section 664 as evaluating how well states and localities provide a free appropriate public education and early intervention services. Assistive technology already sits inside that duty. Each public agency must make assistive technology devices and services available when a child's special education, related services, or supplementary aids and services require them (34 C.F.R. § 300.105). The same regulation ties home use to whether the child "needs access to those devices in order to receive FAPE."

IDEA also defines an assistive technology service to include "a functional evaluation of the child in the child's customary environment" (20 U.S.C. § 1401(2)(A)). The first paper sets the same unit of research, the human functional outcome measured in the person (Puglisi, 2026a). The questions above ask whether a tool class the statute already covers by function is reaching the children it was written for. That question falls within the evaluation of FAPE that section 664 assigns.

Question 6: Data Collection for Future SI Uses

The notice uses Super Intelligence, or SI, for what it formerly called Artificial Intelligence, or AI.

Measure the outcome in the child. A model that scores well on a language test may still fail a blind user whose documents it misdescribes, and a modest model may serve a person well (Puglisi, 2026a). Data built for future SI uses should therefore record what the tool did for the child, not only which tool was used. The first paper lists candidate measures that fit this purpose:

- independent task completion
- task initiation time
- error reduction
- successful communication

- reduced dependence on another person
- completion across interrupted sessions
- user-reported autonomy
- retained decision authority

These measures pair with two established assistive technology outcome instruments, the Psychosocial Impact of Assistive Devices Scale (Jutai & Day, 2002) and QUEST 2.0 (Demers et al., 2002). Both were built for assistive products in general and could be adapted and evaluated for AI uses that pass the functional test (Puglisi, 2026a). Whether current IDEA collections record any of these outcomes for AI tools is a question this comment leaves to NCEE and NCES.

Record the human decision beside the machine output. Under Checkpoint-Based Governance, "There is no AI Governance without human authority and accountability" (Puglisi, n.d.). At each checkpoint, the human governor accepts, modifies, or rejects what the system produced. The same framework separates presence from authority: "Human in the loop means the person is present and participating" (Puglisi, n.d.). An AI-assisted decision about a child's evaluation, eligibility, services, or discipline is not a governed decision until a named human completes it. The record should show who decided, on what evidence, and why (Puglisi, 2026b). Collections designed for SI uses can carry that record next to each machine output, so that any later analysis can tell a governed decision from an automated one.

Protect what the data reveals. Adaptive systems learn sensitive facts about a person's cognitive, psychological, and physical limits, and those facts need protection (Puglisi, 2026a). In the third paper, a diagnosis sits behind the validation the way a physician's certification sits behind a parking placard. The issuer sees the diagnosis at certification and keeps the certification, the functional domain, and the clinician's name and date. The person keeps the diagnosis, and it appears on no face and in no notice (Puglisi, 2026c). The standard requires disclosure to no institution, platform, or reader at any point, and any further disclosure is the person's choice. Data collections built for SI uses face the same choice about where a child's diagnosis travels.

Question 7: Products for Practitioners and Policymakers

The papers offer three products for practitioners and policymakers.

A model accommodation clause for school AI policies. The second paper offers this text (Puglisi, 2026b):

> *Nothing in this AI policy is intended to restrict the use of AI or AI-enabled functionality approved as a reasonable accommodation or other disability-related assistive technology. Requests for such use will be evaluated through the organization's applicable disability accommodation process, subject to legitimate requirements concerning essential functions, security, confidentiality, academic integrity, and fundamental alteration.*

The clause guarantees no request. It guarantees that a request is treated as an accessibility question from the start, and in a school that process is the IEP or the Section 504 plan.

A minimum accommodation layer. The same paper lists eight elements a disability-aware AI policy carries (Puglisi, 2026b):

- an accommodation exception, or a cross-reference to the accommodation process
- individualized review
- documentation proportional to need
- approved secure alternatives
- confidential handling of accommodation information
- an appeal path
- separation of accommodation from ordinary AI enforcement
- human review before any adverse action based on detection or reports

A three-prong screen for IEP teams. The AIDA Functional Test gives teams a plain way to separate assistance from substitution. The first paper works a case the test excludes. A person asks AI to write a report on a subject the person never studied, then submits it unread (Puglisi, 2026a). The use fails because the AI answered the absence of subject knowledge rather than a barrier the disability created. The boundary lets a school grant the tool without granting every use of it.

An existing institutional model shows the practice already works. The Open University states that students who have shared a disability "may use Generative AI as a reasonable adjustment" (The Open University, 2025). The condition is that the use "aligns with the assessment's intended learning outcomes and does not compromise academic integrity."

Question 8: Topics Outside the Six Areas

The gap between eligibility and need. Disability law reaches only people who meet its definitions, and many people with partial or episodic limitations fall short of them (Puglisi, 2026b). A child outside IDEA eligibility may still face a real barrier that general-purpose AI reduces. The second paper states the governing principle as "Same standard, different means," which applies as a matter of fairness before it applies as a matter of law.

Notice and the response record. The third paper proposes a dated notice that names the protection, names the policy or action that conflicts with it, states that AI is the person's disability tool, and explains why. It asks people who serve such notices to report what comes back. The notice has force today against covered employers, funded programs, and public entities, where the sender is a person the statute protects and the need is communicated (Puglisi, 2026c). Force there means a duty to consider and respond under the statute's terms, never a duty to grant the means named. Against platforms and publishers, the notice builds a record. Schools and state agencies sit in that first group, so an evaluation of how they respond to such notices would show whether existing protections reach AI use in practice.

Dissent From Inside the Disability Community

The first paper carries the strongest opposition to its own argument in its own pages. The AI Now Institute report Disability, Bias, and AI argues for moving away from a narrow focus on technology's potential to "solve" and "assist." It calls instead for approaches that account for social context, history, and power (Whittaker et al., 2019). The Autistic Self Advocacy Network

asks organizations not to use generative AI for plain language, on tested evidence of errors that change meaning (Autistic Self Advocacy Network, 2025). Research question 4 above carries that dissent into the evaluation by asking which functional barriers show harm from general-purpose AI.

Standpoint of This Comment

This comment presents governance and policy analysis informed by statutes, regulations, and institutional practice. The legal reading is that of a retired police officer with a Master of Public Administration whose studies included constitutional law, business law, procedural law, courts, and judicial case studies. It is not legal advice, it does not predict the outcome of any dispute, and no attorney reviewed it.

References

Autistic Self Advocacy Network. (2025, July). ASAN says no generative AI in plain language. https://autisticadvocacy.org/2025/07/asan-says-no-generative-ai-in-plain-language

Demers, L., Weiss-Lambrou, R., & Ska, B. (2002). The Quebec User Evaluation of Satisfaction with Assistive Technology (QUEST 2.0): An overview and recent progress. Technology and Disability, 14(3), 101 to 105. https://doi.org/10.3233/TAD-2002-14304

Department of Education, Institute of Education Sciences. (2026, November 13). Request for information (Docket ED-2026-IES-3268, FR Doc. 2026-23155). Federal Register, 91(218), 73218 to 73219. https://www.govinfo.gov/content/pkg/FR-2026-11-13/html/2026-23155.htm

Jutai, J., & Day, H. (2002). Psychosocial Impact of Assistive Devices Scale (PIADS). Technology and Disability, 14(3), 107 to 111. https://doi.org/10.3233/TAD-2002-14305

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. Patterns, 100779. https://doi.org/10.1016/j.patter.2023.100779 (preprint: https://arxiv.org/abs/2304.02819)

Puglisi, B. C. (n.d.). Checkpoint-Based Governance. https://basilpuglisi.com/cbg/

Puglisi, B. C. (2026a). When AI becomes assistive technology: Artificially Intelligent Disability Assistance (AIDA) [Working paper]. Zenodo. https://doi.org/10.5281/zenodo.22908271

Version 2, November 8, 2026: https://doi.org/10.5281/zenodo.24056345

Puglisi, B. C. (2026b). When AI policy becomes an accessibility barrier: ADA exposure from restrictions on Artificially Intelligent Disability Assistance (AIDA) [Working paper]. Zenodo. https://doi.org/10.5281/zenodo.22940224

Puglisi, B. C. (2026c). A call for an AI disability standard: Notice and validation for people who use AI as a disability tool, under protections that already exist [Working paper]. Zenodo. https://doi.org/10.5281/zenodo.24056037

The Open University. (2025, July 15). Generative AI for students. https://about.open.ac.uk/policies-and-reports/policies-and-statements/generative-ai-learning-teaching-and-assessment-ou-0

Whittaker, M., Alper, M., Bennett, C. L., Hendren, S., Kaziunas, L., Mills, M., Morris, M. R., Rankin, J., Rogers, E., Salas, M., & West, S. M. (2019, November). Disability, bias, and AI. AI Now Institute. https://ainowinstitute.org/wp-content/uploads/2023/04/disabilitybiasai-2019.pdf

20 U.S.C. § 1401. Definitions. https://www.law.cornell.edu/uscode/text/20/1401

20 U.S.C. § 1414. Evaluations, eligibility determinations, individualized education programs, and educational placements. https://www.law.cornell.edu/uscode/text/20/1414

34 C.F.R. § 300.105. Assistive technology. https://www.law.cornell.edu/cfr/text/34/300.105

Basil C. Puglisi, MPA
A Human-AI Collaboration

#AIassisted using the HAIA Ecosystem | CC BY-NC-SA 4.0

Free for personal, educational, and noncommercial research use with attribution. Commercial exploitation, paid productization, and enterprise commercialization require separate permission and licensing.

Suggested citation: Puglisi, B. C. (2026). Public Comment on Priority Research Questions for Evaluating IDEA Implementation. basilpuglisi.com.