

When AI Becomes Assistive Technology: Artificially Intelligent Disability Assistance (AIDA)

Artificially Intelligent Disability Assistance (AIDA): How General-Purpose AI Can Help People Work, Communicate, Think, and Create Around Physical and Cognitive Limitations

Basil C. Puglisi, MPA

basilpuglisi.com

Working Paper, September 2026

DOI: <https://doi.org/10.5281/zenodo.22908272>

Keywords: Artificially Intelligent Disability Assistance, AIDA, assistive technology, generative AI, disability, functional barriers, reasonable accommodation, reasonable adjustment, human capability, Checkpoint-Based Governance, Augmented Intelligence

Companion paper: *When AI Policy Becomes an Accessibility Barrier: ADA Exposure from Restrictions on Artificially Intelligent Disability Assistance (AIDA)*

Abstract

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. It may restore access to a capability the person already holds, or it may expand one the disability constrains. This paper names that function Artificially Intelligent Disability Assistance (AIDA). The term generalizes a device name from human-robot interaction research into a functional category. The paper proposes a three-prong functional test for when a particular use of general-purpose AI qualifies as AIDA. 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. It runs seven hypothetical cases through the test. They range from blindness and speech loss to intellectual and developmental disability, post-traumatic stress, and combinations of limitations that sit in gray areas. The paper includes a worked case the test excludes, and a named section for the strongest dissent inside the disability community.

The work rests on a position the author holds about the future of work. That future is Augmented Intelligence, human-AI collaboration that levels the field. A person willing to work and learn can then compete with people who hold greater natural ability or more formal education. AIDA is the strongest case for that equalizing effect. It lets people with disabilities reclaim capabilities a barrier had blocked, and it gives people born with a limitation the support to reach a level of function they could not reach without it.

The literature is converging on this question from several sides, and none of the works located joins every side. Wright (2026) sets criteria for non-generative tools only. Morgan (2026) argues the case for one population under one national law. Survey work, standards, and published principles document the use and state the goal, but they stop short of a test. The contribution here is the synthesis those works point toward: a cross-diagnosis functional test that reaches the generative functions of drafting, restructuring, and summarizing. The companion paper pairs it with an analysis of how institutional AI policies engage disability-law accommodation duties.

Author's Note

How the question arrived. The naming of HAIA, which stands for Human Artificial Intelligence Assistant, is already on the public record (Puglisi, n.d.-a). A shoulder surgery in September 2025 took away the use of my dominant arm. That surgery shifted my workflow from typing to voice, and it produced the name HAIA. That experience is the example of AI's value that brought this subject to my attention. It is not evidence for any claim in this paper, and a temporary impairment does not carry the same legal status as the cases analyzed here.

Author standpoint. This author feels that the industry created a problem with automation that made it the face of the industry, AI Agents and what I call Responsible AI, code is no substitute for humans. The fact that AI has so many positive impacts on humanity should not be dismissed or punished because of "possible" misuse or future issues. As someone who had AI to help when I lost the temporary use of an arm, I could not imagine telling someone with long term or permanent disabilities that they can not use a tool that would help them compensate, compete and have a better life because of the "what ifs". I also believe this holds with United States law, which does not punish people for the tools they use to offset their disabilities, while keeping its existing limits in place: undue hardship, fundamental alteration, essential functions, and legitimate security requirements.

Responsible AI, as the author defines it, is machine checking machine: an AI agent running the pipeline, AI reviewing AI, or a human in the loop who only watches, rarely acts, and barely observes (Puglisi, 2026a).

1. Introduction: The Capability Remains, but the Conventional Route Does Not

A disability can leave research ability, writing knowledge, judgment, and decision authority fully intact while it blocks one of the routes those capabilities normally travel to become work. The blocked route might be sustained physical interaction with a keyboard, or reading dense text for hours. It might be holding many moving parts in working memory, speaking aloud in real time, or staying at a fixed workstation for a full day. When the route closes, observers often see a person who cannot do the work. The accurate description is a person who cannot do the work in the expected way.

General-purpose AI opens other routes. A person who cannot type can dictate and then direct a system to organize what was said. A person who cannot see an image can ask for a description of it. A person whose reading is slowed by dyslexia can hear a dense document read aloud and then question it. A person whose working memory fails under load can hand the load to a system that holds context across sessions. In each of these examples the AI mediates access to the person's own capability, and the capability being claimed stays the person's. It carries the person's thinking past a barrier the disability created, or it helps build a capability the disability had held back.

The central question of this paper follows from that observation:

When a person with a disability uses general-purpose AI specifically to overcome a disability-related functional limitation, is the AI functioning as assistive technology?

The answer offered here is that it can be, and that the answer turns on the use, whatever the tool. The same model that drafts marketing copy for one user can serve as a communication aid for another. The difference lies in the relationship between the person, the barrier, and what the AI does about it.

The paper develops that answer in order. Sections 2 and 3 set out how disability and assistive technology are defined in international and national frameworks. Section 4 introduces Artificially Intelligent Disability Assistance (AIDA) as a functional category, and Section 5 separates it from Augmented Intelligence. Section 6 states a three-prong functional test and a separate governance screen, and works one case the test excludes. Section 7 maps barrier types against AI capabilities. Section 8 runs seven hypothetical cases through the test, covering sensory, communication, learning, developmental, psychological, combined physical, and combined borderline situations. The later sections address the mitigation paradox and the theory the test builds on. They then cover risks and failure modes, dissent from within the disability community, the measurement of functional gain, and the research agenda.

Two notes belong at the start. General-purpose AI here means a commercial system built on a foundation model, such as a chat assistant or a multimodal tool, sold to the public for no particular task. It is not the European Union AI Act's defined term for general-purpose AI models, which carries its own obligations on model providers. Temporary impairments are analyzed on their own terms where they arise. The personal history in the Author's Note does not stand in for the legal status of any case in this paper.

2. Disability Is Not Incapacity, and It Is a Functional Relationship

Disability does not mean the absence of ability. A person may retain expertise, reasoning, creativity, and professional skill while being limited in sitting, standing, manual tasks, concentration, executive function, or communication. The people this paper concerns are limited in the conventional means through which they are expected to act, while the capabilities those means express remain intact.

International frameworks already describe disability this way. The World Health Organization's International Classification of Functioning, Disability and Health treats disability as arising in the interaction between a health condition and the environment. Environmental factors shape it alongside the person. The United Nations Convention on the Rights of Persons with Disabilities (CRPD) builds on that understanding. Its Article 1 describes persons with disabilities as those whose long-term impairments, in interaction with various barriers, may hinder full and equal participation in society. Article 4 sets the general obligations of states parties, and Article 9 addresses accessibility. Article 21 covers freedom of expression and access to information, and Article 27 covers work and employment.

The movement across these frameworks runs from a medical model, which locates the problem in the body, to a social model, which locates it in the environment. A human rights model then treats access as an entitlement, and this paper works inside the second and third. It never asks whether AI fixes a condition, because AI fixes none. It asks whether AI changes the environment through which a person meets a task, and whether that change returns the task to the person's own capability.

3. The Existing Definition of Assistive Technology

Assistive technology is defined by what it does for a person, whoever made it and however it was marketed.

The World Health Organization and UNICEF's *Global Report on Assistive Technology* (2022) treats assistive products broadly. The category includes software and generally available products alongside devices built by disability specialists. That global anchor

matters because it removes the assumption that a mainstream product cannot be assistive.

The CRPD adds the concept of reasonable accommodation in Article 2: “necessary and appropriate modification and adjustments not imposing a disproportionate or undue burden, where needed in a particular case, to ensure to persons with disabilities the enjoyment or exercise on an equal basis with others of all human rights and fundamental freedoms.” The Convention also states that discrimination on the basis of disability includes denial of reasonable accommodation.

National law in the United States carries the same functional logic. The Assistive Technology Act defines an assistive technology device as “any item, piece of equipment, or product system, whether acquired commercially, modified, or customized, that is used to increase, maintain, or improve functional capabilities of individuals with disabilities” (29 U.S.C. § 3002(4)). The Individuals with Disabilities Education Act applies the same test to children and adds the phrase “commercially off the shelf” (20 U.S.C. § 1401(1)). The two definitions are quoted separately here because they differ in wording, and the difference is worth keeping.

Neither definition asks whether the item was designed for disability. Both ask whether it is used to increase, maintain, or improve functional capability. A technology need not be built as a disability product to perform an assistive function, and that principle predates generative AI by decades.

4. Artificially Intelligent Disability Assistance (AIDA)

4.1 Where the term comes from

The acronym AIDA already exists in disability-related research, and this paper builds on that record and keeps its name. Williams, Williams, Moore, and McFarlane (2019) presented a social co-robot for workers with intellectual and developmental disabilities at the 2019 ACM/IEEE Conference on Human-Robot Interaction. Their abstract describes the robot as designed “to uplift and augment the abilities” of those workers. It expands the acronym as “an artificially intelligent disability aide.” This paper cites that work from its abstract, since the body was not available for review.

The same research group then shifted the expansion. Moore and Williams (2020) described AIDA as an “artificially intelligent disability assistant” that used social scaffolding to support workers with intellectual and developmental disabilities. Moore and Williams (2021) carried that wording into the title of a paper on a learning architecture for the system. Across the three papers the term names a particular device serving a particular workplace population. None of the three defines a general category.

The 2020 study also used a method worth reporting neutrally, since its authors simulated disabilities with their participants.

4.2 The generalization

This paper takes four steps away from that origin while keeping its name. It moves from a named device to a functional category. It moves from “aide” and “assistant,” both things, to “assistance,” a function. It moves from an embodied robot to general-purpose software, and from one population in one workplace to any disability-related functional barrier in any setting.

Definition. Artificially Intelligent Disability Assistance (AIDA) is any AI capability a person uses to increase, maintain, or improve a functional capability limited by disability.

The definition deliberately borrows the verbs of the Assistive Technology Act, so that AIDA sits inside the existing legal concept. Its core proposition follows directly:

General-purpose AI can function as assistive technology when a person uses it specifically to increase, maintain, or improve a functional capability limited by disability.

The relationship between three elements decides whether a use qualifies:

Disability-related functional barrier → matched AI capability → restored or expanded human function

The earlier authors and this paper share one commitment, since they positioned their robot as a way to support workers rather than displace them. This paper treats that commitment as the measure of whether AIDA is working.

4.3 Collisions with other uses of the acronym

Other projects use AIDA or near variants, and readers should not confuse them with this category. The EU Horizon 2020 project AIDE developed adaptive multimodal interfaces to assist disabled people in daily activities. The acronym has also named MIT’s Affective Intelligent Driving Agent, an in-car social robot (Williams & Breazeal, 2013). It names Masats’ Artificial Intelligence Door Access system for buses, which uses computer vision to detect passengers with reduced mobility and adjust door and ramp operation (Sustainable Bus, 2025). That system sits closest in purpose to this paper’s category, and it remains a named product, where AIDA here names a functional category. In marketing the acronym has long named a sales funnel model. The searches conducted for this paper did not identify prior use of the exact phrase “Artificially Intelligent Disability Assistance.”

5. AIDA and Augmented Intelligence: Two Standalone Concepts

AIDA is not Augmented Intelligence, and the two concepts stand on their own.

Augmented Intelligence, as the author's framework work defines it, has three parts (Puglisi, 2026b). The human provides the input, the prompt, and the governing structure. AI "performs the execution work to approximately ninety percent completion (a conceptual approximation, not a measured threshold: what it names is the zone where AI delivers maximum value before human judgment must close the remaining gap)." A governed human checkpoint then closes the remaining gap between what the machine produces and what the human is accountable for. The ninety percent figure is theory, as the definition itself states, and this paper treats it that way. Augmented Intelligence extends what a person can do under human governance.

AIDA does something different. It restores access to a capability the person holds, or it expands a capability the disability constrains. The same tool can serve either purpose, and one person can use it for both at once. The statutory verbs control the distinction. What prong 3 asks is whether the use increases, maintains, or improves a functional capability. The functional test in Section 6 sorts a given use into one category or the other. The category depends on what the use does for the person, whichever product performs it.

The bridge between the two concepts is parity. AIDA removes the access barrier so the person can enter augmented practice on equal terms with everyone else. A person who cannot type cannot direct an AI through a keyboard at the speed of the workplace. Voice input closes that gap first, and only then does the person stand where other augmented practitioners stand. The author's measurement work states the boundary between the two concepts precisely: "Accommodation removes barriers to demonstrating capability; it does not erase a real difference in capability" (Puglisi, n.d.-b). AIDA is the accommodation half of that sentence. Augmented Intelligence is the capability half.

The equalizer. The two concepts meet in one larger claim. The author holds that the future belongs to Augmented Intelligence, and that human-AI collaboration narrows the distance that intelligence, schooling, and circumstance once fixed between people. A worker without an advanced degree who is willing to work and learn can direct AI well enough to compete with credentialed peers, and a person of ordinary ability can hold ground against people of exceptional ability. That is an argued proposition about the future of work, and this paper treats it as argument. AIDA is its clearest instance. For a person with a disability, the same collaboration does more than narrow a gap in advantage. It returns a capability the barrier had taken, or it gives a person born with a limitation the support to reach a level of function that was never available before.

Two clarifications prevent confusion about how the terms relate. The original AIDA authors used "augment" in its everyday sense, meaning to add to, which differs from

Augmented Intelligence as a term of art. And the division of labor in the Augmented Intelligence definition describes a zone of value that remains to be measured.

6. A Functional Test for AIDA

The test comes before the cases so that each case can run through it, and it has two parts. The AIDA Functional Test, three prongs, decides whether a use of AI performs an assistive function at all. The Governance and Effectiveness Screen, two further prongs, asks a separate question about consequence and accountability once a use qualifies. The screen is the paper's own addition, drawn from Checkpoint-Based Governance, and it does not decide assistive status.

AIDA is governance-neutral. The functional test asks what the AI does for the person and leaves the governance model aside. A tool with a human checkpoint qualifies. So does a product that validates its own output, and so does an agent that carries out a sequence of steps on the person's behalf. For some people the barrier is precisely the inability to perform the steps, and delegation to an agent is itself the accommodation. Nothing in prongs 1 to 3 turns on the governance model, and no person loses assistive status because the tool they can use runs without their approval at each step.

The three-prong functional test and the separate screen




Figure 1. The three-prong AIDA Functional Test decides assistive status. The separate Governance and Effectiveness Screen asks who answers for consequential use and whether the use delivers a benefit.

6.1 The AIDA Functional Test

1. **A disability creates an identifiable functional barrier.** The barrier must be nameable in functional terms, such as reading dense text, sustaining a keyboard session, holding sequence in working memory, or speaking in real time. A diagnosis alone is too broad to meet the prong.

2. **A particular AI capability directly reduces or bypasses that barrier.** The match must be specific. Voice input answers a typing barrier, visual description answers a seeing barrier, and task decomposition answers a sequencing barrier; a general claim that AI “helps” does not satisfy the prong.
3. **The intervention increases, maintains, or improves the person’s functional capability.** This prong borrows the statutory verbs directly. It covers restoring access to a capability the person holds, sustaining a capability that would otherwise erode under load, and expanding a capability the disability had constrained.

6.2 The Governance and Effectiveness Screen

4. **Someone holds authority over consequential AI-mediated activity, and it is the person whenever the person wants it.** The person’s intention starts the activity, and the authority a checkpoint carries is the authority to approve, override, modify, or escalate what the system produces (Puglisi, n.d.-c). Where the person delegates a step to an agent or to a tool that runs unattended, accountability for that step moves to whoever deployed it. The question becomes who answers, and the use keeps its assistive status. The author’s measurement work names the underlying competence as the ability to “direct, challenge, verify, and own the work done with AI” (Puglisi, n.d.-b). This prong applies the human authority requirement of Checkpoint-Based Governance, under which “there is no AI Governance without human authority and accountability” (Puglisi, n.d.-c). That authority rests on no credential. In the framework’s own terms, the requirement “does not depend on prior practice, credentials, or seniority” (Puglisi, n.d.-c), and a disabled person’s authority over the person’s own assistive use stands on the same ground.
5. **The intervention produces, or is reasonably capable of producing, an observable functional benefit.** The benefit need not be proven in advance for every person, since a new accommodation is often tried before it is known to work. It must be the kind of benefit a person or an observer could see.

The screen is not a condition of assistive status, and a use that fails it is still assistive. Automatic captioning in a lecture hall passes the functional test whether or not the person governs it, and so does an agent that books an appointment for someone who cannot make the call. What the screen asks is a question every AI user faces, disabled or not: when an output carries consequences, who answers for it. That question belongs to the activity and to the institution around it, and it is why the screen sits beside the definition. Checkpoint-Based Governance draws the same line in its distinction between a human who is merely present and a human who holds authority and bears accountability (Puglisi, n.d.-c).

The full chain reads:

Human capability → disability-related barrier → AI intervention → restored or expanded functional access → human action

The functional chain




Figure 2. *The functional chain from human capability to human action.*

6.3 A worked case the test excludes

A person with a documented disability asks a general-purpose AI to write a professional report on a subject the person has not studied, then submits the report unread. The use fails the functional test at prong 2, because the AI capability does not answer the disability-related barrier; it answers the absence of subject knowledge, which no disability created. It fails prong 3, because no capability of the person is increased, maintained, or improved. The governance screen never comes into play, since the use is excluded before the screen is reached. The exclusion rests on the functional test itself and passes no judgment on the person or the disability.

The exclusion matters for two reasons. It shows the category has a boundary that ordinary institutions can apply. It also answers in advance the charge that any AI use by a disabled person will be called an accommodation. The test asks what the AI did for the person's own capability, and a use that replaced that capability does not qualify.

7. The Barrier and Capability Matrix

Diagnosis-by-diagnosis coverage cannot scale, because the same diagnosis produces different barriers in different people and different diagnoses produce the same barrier. The matrix below organizes the question by barrier type instead. It supports the claim that AIDA reaches every category of functional barrier, including limitation that is partial, episodic, or short of a formal disability determination. The emerging and contested rows mark where direct study remains thin, and the claim extends across them. The matrix does not claim that AI helps every person within a category, since the same tool can serve one person and fail another.

The same breadth carries a moral point. A person whose disability is partial, or whose limitation falls short of a legal threshold, still meets the barrier the matrix names. Taking away a tool that reduces that barrier raises a moral and ethical question whether

or not any law requires the tool. This paper holds that such a removal needs a stated reason tied to a real risk.

Barrier type	AI capabilities that can match it	Evidence status
Physical input and mobility	Voice input and output, drafting from dictation, persistent context across sessions	Documented in practice; speech recognition software was evaluated as an assistive device before generative AI (DeRosier & Farber, 2005)
Vision	Visual description, document reading, question answering about images	Documented with blind users, including verification and hallucination problems (Adnin & Das, 2024; Gonzalez Penuela et al., 2026)
Hearing and Deaf communication	Transcription and captioning, text-based summaries of spoken exchanges	Documented in workplace use (Shock & Silverman, 2026); caption accuracy varies
Speech and communication	Text to natural speech, phrase suggestion, rephrasing for spoken exchange	Documented with AAC users (Valencia et al., 2023)
Reading and written expression	Text to speech, summarization, restructuring of the person's own drafts	Documented with disabled students (Zhao et al., 2025; Zhao et al., 2026)
Executive function and working memory	Task decomposition, sequencing, persistent context, prompted recall	Documented in cross-disability research (Johnson et al., 2026)
Attention and task continuity	Continuity across interrupted sessions, adaptive pacing	Emerging; limited direct study
Stress tolerance and interpersonal load	Asynchronous preparation, drafting before difficult communication	Emerging; limited direct study
Autism and sensory regulation	Rephrasing tone, preparing for social exchanges	Documented in workplace communication research (Jang et al., 2024)
Acquired brain injury	Sequencing, memory support, plain-language rephrasing	Emerging; limited direct study
Chronic illness, pain, and fatigue	Shorter work cycles, voice input, continuity across rest periods	Emerging; limited direct study
Episodic capacity	Continuity across fluctuating availability	Emerging; limited direct study
Plain-language access	Rephrasing into plain language	Documented in use, with meaning-changing errors reported in testing. See Sections 8 and 12
Multiple disabilities	Combinations of the above	See the combined cases in Section 8

Cells marked emerging carry a mechanism and little direct study, and they belong on the research agenda in Section 15 until study catches up. The plain-language cell carries a live dispute over institutional use, which Sections 8 and 12 set out.

8. Seven Cases

All seven cases are hypothetical, and none is linked to the author. Each runs through the functional test and the governance screen in Section 6. Four sit in Tier A, where both the disability and the functional link are clear. Three sit in Tier B, the gray area where one or the other invites dispute.

8.1 Tier A: clear-cut physical cases

Blindness. A blind professional receives documents with embedded charts and images that screen readers cannot interpret. The professional uses AI visual description to convert the visual content into language, then questions the description until the content is clear. The barrier is identifiable (prong 1), the capability matches it directly (prong 2), and the professional’s existing analytical capability is restored to the full document (prong 3). The professional directs every description and can reject any of them (prong 4). The benefit shows in documents read that could not be read before (prong 5). The evidence supports the benefit and names its cost. In the American Foundation for the Blind’s national survey, 79% of blind and low vision participants used AI visual description, and 51% of those users did so at least daily. Twenty-one percent reported an AI description error that hurt them, against 9% of sighted users, because a blind user may have no independent way to check visual details against the source image (Hanuschock & Silverman, 2026). In interviews with 19 blind users of generative AI, Adnin and Das (2024) documented real use alongside accessibility, hallucination, and verification problems. A diary study with blind and low vision people found the same pattern in daily use (Gonzalez Penuela et al., 2026). The case passes the test, and the verification burden stays with it.




Figure 3. AI visual description among blind and low vision participants, and the rate of harmful description errors compared with sighted users. Source: Hanuschock and Silverman (2026).

Speech loss. A person who has lost speech communicates through typed text converted to natural speech. The barrier in real-time conversation is speed, since

composing each message takes longer than the exchange allows. AI-mediated communication reduces that barrier through phrase suggestions, rephrasing for spoken exchange, and completion of short replies. This case sits inside the long augmentative and alternative communication (AAC) tradition and is the only one in this paper built on communication as the barrier. Valencia et al. (2023) tested live language-model suggestions with 12 AAC users. The participants reported savings of time and of physical and cognitive effort, and they insisted that the suggested phrases reflect their own communication style. That insistence is prong 4 in the users' own words: the value of the tool depended on the person keeping authorship of what was said.

8.2 Tier A: clear-cut mental and developmental cases

Dyslexia. A student with dyslexia reads dense assigned texts slowly and loses meaning under time pressure, while comprehension of the same material heard aloud is strong. The student uses text to speech, then summarization and question answering to check understanding, and uses AI to restructure the student's own drafts. The accommodation lineage here is long, running from spell check to text to speech to generative AI, and the functional logic has not changed across it. Zhao, Cox, and Chen (2025) surveyed 124 disabled students and found generative AI used to overcome academic writing barriers, alongside concerns about accuracy, academic integrity, and subscription cost. The case passes the functional test and the governance screen where the student's own ideas remain the content being organized.

Intellectual and developmental disability. A worker with an intellectual disability performs a multi-step job reliably once the sequence is clear, and loses the sequence when steps change. AI supports task scaffolding, step sequencing, plain-language rephrasing of instructions, and prompted recall. This is the population of the original AIDA papers, so the case carries the lineage of Section 4 into general-purpose software. Here AIDA often expands capability, which is why prong 3 reads "increase, maintain, or improve."

The evidence in this case is mixed, and the case carries the mix. Randall, Drew, Gilman, and Dixon (2025) interviewed 25 people, 22 adults with intellectual and developmental disabilities and three job coaches, and found assistive technology used successfully and independently at home. Use in workplaces was extremely limited. The study documents a gap between settings, with workplace barriers among the constraints participants reported. Coffey, Krahn, Hanley, and Neely (2026) found that stories generated by large language models portrayed people with intellectual disabilities as younger and more dependent. The same stories cast them as more in need of help than people without disabilities.

The plain-language function carries a direct dissent. The Autistic Self Advocacy Network tested multiple models and found errors that changed meaning. It asks organizations

not to use generative AI for plain-language translation (Autistic Self Advocacy Network, 2025). Rodgers (n.d.), writing for the University of Minnesota's Institute on Community Integration, argues that AI used with guardrails and human expertise can expand plain-language access. Rodgers also reports people with intellectual and developmental disabilities already using it.

The worker is the qualified person here, and the use is the worker's own. The test asks the same question it asks in every other case, whether the use increases, maintains, or improves the worker's functional capability. Where the plain language carries the meaning of the original instructions, the case passes. Where a meaning-changing error goes unnoticed and the worker acts on the wrong instruction, that output delivers no prong 3 benefit, and a use that changes meaning systematically fails the prong. The ASAN finding is about organizations generating plain language on people's behalf, which is a different arrangement from a worker running the tool on instructions the worker was handed.

8.3 Tier B: gray-area cases

Post-traumatic stress disorder. No profession and no cause are named, because the case turns only on functional barriers. A person with post-traumatic stress disorder (PTSD) has difficulty concentrating under stress, sleeps irregularly, loses task continuity when interrupted, and finds some real-time interpersonal exchanges overwhelming. AI offsets these barriers through asynchronous communication, structured organization of the person's own thoughts, flexible pacing, preparation before difficult communication, and continuity across interrupted work. The boundary with clinical treatment must hold. Organizing communication during psychological stress is assistive; offering therapy is a different function, and this paper makes no claim for it.

This is a borderline case, and its qualification as a disability is part of the debate. The EEOC regulation names post-traumatic stress disorder among the impairments that should easily be found to substantially limit brain function, which weighs strongly toward coverage (29 C.F.R. § 1630.2(j)(3)(iii)). The determination still rests on an individualized assessment of the particular person (29 C.F.R. § 1630.2(j)(1)(iv)), and an institution may contest qualification in a particular case. The functional link draws interpretation as well: whether a given AI use answers a PTSD-related barrier, as prong 2 requires, or answers something else. Both questions are decided case by case, which is why the case sits in Tier B.

Combined physical limitation. A person has a combination of physical injuries, none named here, which alone may not amount to a substantial limitation. Together they limit walking, standing, sitting, typing, or writing across the eight hours a day that many conventional full-time roles expect. The person's capability for the work itself is intact. What fails is the capability for the conventional structure of the work, and AI

supplies a different means. Voice input replaces typing, and the person can change position without losing the thread of the task, work in shorter cycles, and work away from a fixed desk.

The case needs no special rule about aggregating impairments. It stands on the combined functional effect on activities that are themselves recognized, such as walking, standing, sitting, and performing manual tasks. Its gray quality comes from the individual injuries, each of which invites dispute on its own, while the combination does not.

Combined borderline person. The final case joins the previous two. A person lives with PTSD and the combined physical limitation together and uses AI to offset both at once. Voice input and position changes answer the physical barriers, while asynchronous preparation and continuity answer the psychological ones. The functional test is built for exactly this person, because the person's situation defeats any diagnosis-by-diagnosis approach. Where the functional test holds for each barrier, the value question is answered. The protected-status question, which the law answers separately, belongs to the companion paper.

8.4 Why Tier B is gray

Tier B cases are gray for two different reasons. In the PTSD case, qualification as a disability is open to debate even though the regulation weighs toward coverage, and the functional link to AI draws interpretation as well. In the combined physical case, the individual impairments invite dispute that the combination does not. The functional test resolves the value question in both, because it asks what the AI does for the person, whatever label the person carries.

9. The Mitigation Paradox

Success with technology does not erase disability. When a person performs well with assistive technology, observers may conclude that the person was never limited. The reverse is closer to the truth. Performance enabled by technology can coexist with real disability, and improved performance may be the clearest evidence that the intervention works.

United States law addresses this directly. The Americans with Disabilities Act, as amended, requires that substantial limitation be assessed without regard to the ameliorative effects of mitigating measures. The statute lists the "use of assistive technology" among those measures (42 U.S.C. § 12102(4)(E)(i)). The implementing regulation carries the same rule (29 C.F.R. § 1630.2(j)(1)(vi)). The legal system already refuses to treat a working accommodation as proof that no disability exists. The same reasoning applies when the accommodation is general-purpose AI.

10. From Accommodation to Capability Restoration and Expansion

Traditional assistive technology usually serves one predefined function. A ramp changes the building, and a screen reader carries an interface to the user, while modern assistive technology already includes adaptive and predictive software. General-purpose AI extends that trajectory, because a single adaptive system can mediate several forms of access across tasks, modalities, and contexts. It turns typing into speech, dense text into conversation, large assignments into sequences, visual information into language, and long records into summaries. The same system reshapes its output for a blind user in one session and a dyslexic user in the next.

10.1 The theory this builds on

AIDA does not replace the established models of rehabilitation science. It puts them into operation for general-purpose AI.

- **The Human Activity Assistive Technology (HAAT) model**, developed by Albert Cook, Sue Hussey, and Janice Polgar, places assistive technology in the relationship between the human user, the activity, and the context (Polgar et al., 2025). The functional test asks the same questions of an AI use: which person, which activity, which barrier, which context.
- **Matching Person and Technology** treats the fit between the person and the tool as the determinant of success (Scherer & Craddock, 2002). The test's second prong is a matching requirement.
- **The capability approach** supplies the vocabulary the paper already uses, what a person is actually able to do and to be. Nussbaum (2006) extends that approach to disability, Burchardt (2004) and Mitra (2006) treat disability as a deprivation of capability, and Bickenbach (2014) examines how the approach fits the ICF.

The tools themselves should meet recognized accessibility standards, since an inaccessible assistant defeats its own purpose. WCAG 2.2 governs web content. EN 301 549 is the European standard for accessible information and communication technology. Its version 4.1.1 was adopted on August 24, 2026 and published on September 2, 2026 to support the European Accessibility Act (Davis Wright Tremaine, 2026). Version 3.2.1 remains the standard cited in the Official Journal, which carries the presumption of conformity, until version 4.1.1 is cited there (Deque, 2026).

10.2 Frameworks already in force

Two national instruments already address this ground, and both should be read precisely.

- **Canada.** CAN-ASC-6.2:2025, *Accessible and Equitable Artificial Intelligence Systems*, is a voluntary National Standard of Canada published December 3, 2025. It was developed by a committee made up primarily of people with disabilities.
- **Australia.** The National Disability Insurance Scheme framework for AI-enabled assistive technology was developed by CSIRO's Australian eHealth Research Centre around six principles: user experience, value, quality, safety, privacy and security, and human rights. Silvera-Tawil et al. (2025) report the same work as the AI-enabled AT Framework, so the two form a single lineage.

Both instruments evaluate AI systems and AI-enabled assistive products for development, procurement, and funding decisions. Neither sets out a use-level test for whether a general-purpose tool, in a particular use, is functioning as assistive technology for a particular person. That difference is the space the functional test occupies.

10.3 The research literature

The empirical record on disabled people's use of generative AI is growing quickly. Zhao, Cox, and Chen (2025) and Zhao, Chen, and Cox (2026) documented how disabled students use generative AI in academic writing, including the affordances the tools offer and the costs they impose. Adnin and Das (2024) interviewed blind users, and Gonzalez Penuela et al. (2026) followed blind and low vision people through a diary study. Johnson, Lewis, Mankoff, and Banner (2026) ran seven cross-disability focus groups with 20 participants and found generative AI supporting autonomy, efficiency, and communication while imposing costs the authors call accessibility taxes. Glazko et al. (2023) conducted a three-month autoethnography of generative AI for accessibility.

Silverman et al. (2026), for the American Foundation for the Blind, surveyed 1,735 United States adults, 1,070 of them with disabilities, on the benefits, risks, and aspirations of AI. A secondary analysis of the 1,374 workers in that sample found that among workers with disabilities, 32% used AI for visual description and 25% for captioning on the job. The authors describe that use as AI functioning as assistive technology (Shock & Silverman, 2026). A second analysis of the same survey examines visual description, captions, and augmentative and alternative communication directly as AI assistive technology (Hanuschock & Silverman, 2026). Hadar Souval et al. (2025) frame generative AI as a cognitive copilot for people with cognitive disabilities, offering personalized assistance in daily tasks and social interaction, and as a social mirror that can amplify bias.




Figure 4. AI use at work in a sample of 1,374 workers. Source: Shock and Silverman (2026).

Institutions have begun to act on the same understanding. The Open University permits students who have shared a disability to use generative AI as a reasonable adjustment in any assessment category. The use must not compromise the intended learning outcomes or academic integrity. It does not require that use to be acknowledged (The Open University, 2025). Defra, the United Kingdom’s Department for Environment, Food and Rural Affairs, made full Microsoft Copilot licences available to staff as a reasonable adjustment. It did so after trials showed benefits for neurodivergent and disabled colleagues (Ede, 2026).

11. What AIDA Is Not, and Where It Fails

The category has edges, and the paper names them.

Disability plus AI use does not equal assistive technology. The failing case in Section 6.3 shows where the line falls, because the use replaced the person’s capability. AI performing a person’s substantive professional thinking differs from AI giving access to that person’s own thinking. Offering therapy differs from organizing communication during psychological stress. Convenience differs from accommodation, although a single feature can be convenience for one person and accommodation for another.

The fraud objection. Critics predict that everyone will claim a disability once AI accommodations are possible. Existing accommodation processes already provide ways to establish disability and accommodation need where documentation is lawfully appropriate. Outside formal processes, the functional test itself is the limit, because it asks what the AI did for the person’s own capability.

AIDA also carries real risks:

- **Substitution.** Institutions may use inexpensive AI to cut human accommodations and services, or to replace funded specialist assistive technology. Cost-driven substitution can fail as an effective accommodation even when the substitute can perform an assistive function.
- **Overreliance.** Scaffolding can displace a capability the person retains and wants to keep using.
- **Accessibility taxes.** Johnson et al. (2026) name financial, disclosure, labor, prompt-engineering, and verification costs that fall hardest on disabled users who face multiple forms of marginalization. Disabled students already report subscription cost as a barrier (Zhao et al., 2025).
- **Inaccessible AI.** The tool itself can be inaccessible in its interface, its output formats, its authentication steps, or its price.
- **Verification burden.** A fabricated output lands hardest on the person least able to check it.
- **Identity misrepresentation.** Generated output can misrepresent or erase a user's identity, as Johnson et al. (2026) report.
- **Abandonment and payment.** Assistive technology abandonment research explains why tools fall out of use. Phillips and Zhao (1993) found 29.3% of devices abandoned among 227 adults, with the user's input in selection among the predictors. Gilly (2026a) finds that the main abandonment taxonomies have no category for a cost the supplier changes after the user has adopted the tool. Funding responsibility varies by jurisdiction, program, employment relationship, and accommodation regime. This review identified no single general rule for who pays for a general-purpose AI subscription used as assistive technology, whether vocational rehabilitation, public insurance, an employer, or the person.
- **Vendor dependency.** A model retirement, a policy change, or a change in price or capability can remove an accommodation overnight. Gilly (2026a) names that harm accommodation repricing. The author's framework treats provider plurality as a form of accessibility resilience.
- **Privacy.** Adaptive systems learn sensitive facts about a person's cognitive, psychological, and physical limits, and those facts need protection.

12. Dissent From the Disability Community

This section carries the strongest opposition to the paper's argument in the paper's own pages.

The first group of critiques challenges the framing itself. Whittaker et al. (2019), in the AI Now Institute report *Disability, Bias, and AI*, argue for moving away from a narrow focus on technology's potential to "solve" and "assist" and toward approaches that account for social context, history, and power. This is the most direct challenge to an

assistive framing, and it deserves the first reply. Technoableism scholarship, developed in Shew's *Against Technoableism* (2023), critiques the framing that aims technology at disabled people when the fix belongs in the environments that exclude them. El Morr, Kundi, Mobeen, Taleghani, El-Lahib, and Gorman (2024) reviewed 45 studies and found a narrow medical model and an ableist perspective prevalent across AI and disability research. van Toorn, Scully, and Gendera (2025) examine disability justice and relational ethics in AI. Mladenov (2025) argues that AI can reinforce exclusion through the overvaluing of self-sufficiency, algorithmic injustice, and techno-fetishism, and that it enables independent living only when aligned with disability rights principles.

The second group reports harm and mistrust in use. The Autistic Self Advocacy Network (2025) asks organizations not to use generative AI for plain language, on tested evidence of errors that change meaning. Coffey, Krahn, Hanley, and Neely (2026) found implicit bias against people with intellectual disabilities in stories generated by large language models. Glazko et al. (2023) report concerns about verifiability, training data, ableism, and false promises alongside real accessibility uses. Johnson et al. (2026) found participants using generative AI out of necessity while mistrusting it, a balance between access and risk.

One voice from the same field pulls the other way. Rodgers (n.d.), writing for the Institute on Community Integration, argues that AI used with guardrails and human expertise can expand plain-language access, with disabled people designing and testing the tools.

The paper's reply. AIDA is defined by the person's own functional gain, which is the same standard these critiques apply. Whittaker et al. warn against treating technology as the solution to disability. AIDA treats technology as a route around an environmental barrier, used by or for the person to reduce or bypass it, and the social model in Section 2 locates that barrier where these critics do. Where evidence shows harm, as in the plain-language dispute, the test asks what the use did for the person, and a use that produced a wrong result for that person fails prong 3. The test does not ask whether AI is good for disabled people as a class. It asks whether a particular use served a particular person's own capability, and several of these critiques would accept that question as the right one.

13. Human Agency as the Measure of Success

AIDA succeeds when the person becomes more capable of acting. That measure says nothing about how much the system does on its own. A person who hands an agent a whole task is as well served as one who checks every line, provided the delegation

reduces or bypasses the disability-related barrier and leaves the substantive capability being claimed with the person.

Four conditions follow from the governance screen, and they govern consequential activity, leaving assistive status to the functional test. The person's intention starts the activity. The person can inspect, change, reject, or stop what the system produces, whenever the person wants to. Decisions that carry consequences have a named human who answers for them. Uncertainty stays visible to the person, however fluent the output.

Checkpoint-Based Governance states the governing distinction in its canonical text. "Human in the loop means the person is present and participating. It does not require that the person holds authority or answers for the outcome" (Puglisi, n.d.-c). The same text names the failure the distinction guards against: "The failure mode in AI governance is not the absence of humans. It is presence without accountability" (Puglisi, n.d.-c). An assistive use in which the person is merely present, while the system decides and nobody answers for the result, fails prong 4. It remains assistive, and what it lacks is an accountable human, which the person's own disability never supplied and never should have been asked to. The Defra employee quoted in Section 10.3 described the same condition in plain terms. The employee reviews every output and calls the tool a thinking partner rather than a decision maker (Ede, 2026).

14. Measuring Functional Gain in the Person

The unit of research is the human functional outcome, measured in the person. 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.

Candidate measures include:

- independent task completion
- reduction in physical effort
- task initiation time
- sustained productive periods
- error reduction
- successful communication
- reduced dependence on another person
- cognitive effort
- the ability to change working position without losing the task
- completion across interrupted sessions
- user-reported autonomy
- retained decision authority

These measures pair with established assistive technology outcome instruments. The Psychosocial Impact of Assistive Devices Scale uses 26 items across adaptability, competence, and self-esteem (Jutai & Day, 2002). QUEST 2.0 uses 12 items to measure satisfaction with the device and with services (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.

15. The Contribution and the Research Agenda

15.1 The contribution, stated precisely

Existing assistive technology theory already treats function, person, and context as central. Emerging studies document generative AI use by specific disabled populations. A voluntary Canadian national standard and an Australian funding framework already address accessible AI and the evaluation of AI-enabled assistive products. Institutions already grant general-purpose AI as a reasonable adjustment, as the Open University and Defra show.

This paper's contribution is integration and operationalization. It offers a cross-diagnosis functional test for when a particular use of general-purpose AI functions as assistive technology. The companion paper pairs that test with an analysis of how institutional AI policies engage disability-law accommodation duties.

The nearest prior works each approach this synthesis from one side.

Morgan (2026) argues under United Kingdom law that regulated use of large language models may be a reasonable adjustment where a neurodivergent student's barrier lies in written expression rather than reasoning. That is the principle of same standard, different means, applied to one population in one setting.

Wright (2026) proposes four operational criteria (fidelity, non-augmentation, traceability, and attestation) to separate non-generative AI tools such as optical character recognition, voice-to-text transcription, and handwriting recognition from academic misconduct under institutional restrictions on generative AI. It also argues that detector output is not sufficient for a misconduct finding. Its scope stops at non-generative tools, and the generative functions this paper's cognitive cases rely on (drafting, restructuring, and summarizing) are the ones it excludes. Its criteria can be read as the non-generative half of the functional test offered here.

Brand, Fischer Mogensen, and Engelbrecht (2026) treat generative AI as reasonable accommodation in employment across disability types, without a functional test or a policy taxonomy.

Eberhardt (2026) analyzes generative AI writing tools as cognitive scaffolding for documented memory impairment under the Americans with Disabilities Act and Section 504 of the Rehabilitation Act, one impairment in one setting.

Giraldo and Sacchi (2026) review 16 evaluation approaches for AI-enabled assistive technology and propose assistive validity, the extent to which a system's performance becomes accessible, controllable, and meaningful support for a particular person. Their unit of analysis is the evaluation instrument, while this paper asks when a particular general-purpose AI use functions as assistive technology.

Hanuschock and Silverman (2026), for the American Foundation for the Blind, open from the Assistive Technology Act definition and argue that AI suits assistive use because it converts information across modalities. They analyze visual description, captions, and augmentative and alternative communication in a national survey. The report covers three functions and does not offer a cross-diagnosis test or a policy analysis.

The American Foundation for the Blind (2025) states the principle that students and employees with disabilities should be able to use AI as assistive technology, and that schools and employers should consult the disability community when writing policies on AI as a reasonable accommodation. Its survey work documents the use (Silverman et al., 2026; Shock & Silverman, 2026).

Gilly (2026a) analyzes accommodation repricing and expressly sets aside the disability-law question as a separate one. Gilly (2026b) treats the accommodation duty in disability law as an affirmative obligation distinct from the analysis of discrimination by intent or impact.

Knobel (2026) argues that workplace AI can erode reasonable accommodation for neurodivergent workers in New Zealand.

Mladenov (2025) maps AI-mediated barriers and enablers of independent living, and Della Penna, Buzzi, and Leporini (2026) frame generative AI as assistive technology for web interaction.




Figure 5. How many of 16 published evaluation approaches for AI-enabled assistive technology explicitly operationalize each dimension. Source: Giraldo and Sacchi (2026).

The claim, dated and scoped. Searches concluded September 21, 2026, covered web-indexed scholarly, standards, government, and practitioner sources, and a five-platform cross-AI review followed. Follow-up checks on September 22, 2026 added the two works an outside review surfaced. Neither located a work combining a cross-diagnosis functional test that reaches generative functions with a multi-type, multi-setting analysis of how institutional AI policies (bans, process restrictions, disclosure mandates, distribution penalties, discipline, and detection) engage disability-law accommodation duties. The field is converging on this join, and the works listed above approach it from one side. Appendix A sets out the search method. The contribution here is the synthesis those works point toward, and it claims no priority over the broader idea that AI can function assistively.

Providing another route from human intention to human action.

The hypothesis is not that AI benefits every disabled person. It is that general-purpose AI creates a class of adaptable assistive capabilities across otherwise unrelated barriers.

15.2 Research agenda

The questions that follow are open:

- When does general-purpose AI meet the functional definition, and how consistently can different evaluators apply the test to the same use?
- Which barriers show benefit, and which show harm?
- How do temporary, episodic, long-term, and permanent disabilities differ in what AIDA offers?
- Does AI change employment participation for disabled people, and do the gains persist?
- Where does dependence emerge, and which tasks should stay human-only?

- Should general-purpose AI enter formal accommodation processes, and on what documentation?
- How do accessibility taxes distribute across disabled users?
- How does AIDA operate outside work and school, in healthcare, public benefits applications, voting, legal proceedings, and independent living?
- What changes where subscription cost and connectivity limit access, including in the Global South?

16. Jurisdictional Examples

Each example below is an illustration, and none is the foundation of the argument.

- **United Kingdom.** The Equality Act 2010 imposes a duty to make reasonable adjustments. *University of Bristol v Abraham* separates competence standards from the methods used to assess them, and Morgan (2026) applies that distinction to large language models.
- **Canada.** The Accessible Canada Act, and CAN-ASC-6.2:2025, a voluntary national standard on accessible and equitable AI.
- **Australia.** The Disability Discrimination Act 1992, and the NDIS framework for AI-enabled assistive technology.
- **New Zealand.** Knobel (2026) examines the gap between non-binding AI guidance and reasonable accommodation for neurodivergent workers.
- **United States.** The Assistive Technology Act, the Individuals with Disabilities Education Act, and the mitigating-measures rule of the Americans with Disabilities Act. The companion paper treats these in detail.

17. Implications for Employment and Human Capability

AI may permit some people who possess valuable capabilities but cannot reliably operate within conventional work structures to participate in forms of work that were previously inaccessible or unsustainable.

Capacity for work and capacity for conventional full-time work are not the same thing. The combined physical case in Section 8 shows the difference. A person whose body cannot sustain eight hours at a desk may sustain eight hours of productive work. Voice input, position changes, and shorter cycles take the place of the desk. The capability was never the missing element; the route was.

The evidence on workplace access points the same way. Adults with intellectual and developmental disabilities use assistive technology successfully at home and rarely at work (Randall et al., 2025). Workers with disabilities report assistive uses of AI that workplace policy does not allow (Shock & Silverman, 2026). The same reasoning places

AIDA inside the author's larger claim about work. If human-AI collaboration is what lets a willing learner compete with people who hold more natural or credentialed advantage, then people with disabilities stand to gain the most from it. The barrier they face is the kind the collaboration is best placed to remove. How institutions respond, including the policies that restrict or penalize this use, is the subject of the companion paper.

18. Conclusion: The Human Capability Was There

The disability does not create the capability, and AI does not create it either. The disability blocks one route, and AI provides another; for some people it also helps build what the barrier had held back. Seven cases share one mechanism, and the functional test states when that mechanism is working and when it is not.

Seen from the author's view of the future, AIDA is the equalizer at its most concrete. Augmented Intelligence levels the field for anyone willing to work and learn, and for a person with a disability it can return access to the field itself.

The test keeps two commitments at once. It protects the person who uses AI to reach capability the person already holds, and it refuses the label of assistance to any use that answers something other than a disability-related barrier. It asks nothing about how the tool is built or governed, because a person who needs a task carried out is not served by a rule that only counts tools they supervise step by step.

The most important measure of assistive AI may not be what the machine can do for a person, but what the person becomes able to do because the barrier between capability and action has been reduced.

Appendix A. Novelty Search Method

The contribution statement in Section 15.1 rests on a dated search, and this appendix records how it ran so that another researcher can repeat it.

Dates. The searches concluded September 21, 2026. A follow-up live audit on the same date checked every single-platform lead the review had raised, and follow-up checks on September 22, 2026 retrieved and verified the two works outside reviews then identified.

Source classes. The search covered web-indexed scholarly literature (journal articles, conference proceedings, SSRN, Zenodo, and arXiv), national and international standards, government guidance, institutional policies, and practitioner publications.

Query families. Queries combined three groups of terms: generative AI, large language models, or AI tools; assistive technology, reasonable accommodation, or reasonable adjustment; and disability, neurodivergence, or a named impairment. Setting terms narrowed each run to employment, higher education, academic integrity, and publishing. Named-instrument queries covered CAN-ASC-6.2, the NDIS framework for AI-enabled assistive technology, and EN 301 549. Outside reviews after the search identified two works it had missed, the American Foundation for the Blind's *Innovation for Access* (August 2026) and Giraldo and Sacchi's review of evaluation approaches (September 2026). Both are now named in Section 15.1.

Cross-platform review. Five AI platforms ran the same search brief in parallel: ChatGPT, Gemini, Grok, Meta AI, and Kimi through Perplexity. A human arbiter compared their returns, and every work named in Section 15.1 was then retrieved from its primary record, meaning the journal page, DOI record, repository record, or publisher page. A work raised by one platform and not retrievable from a primary record was left out.

Limits. The search reached web-indexed sources only. Paywalled full texts were read where available and otherwise assessed from their abstracts, and non-English literature was not searched systematically.

References

Primary sources

- Adnin, R., & Das, M. (2024). “I look at it as the king of knowledge”: How blind people use and understand generative AI tools. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '24)*. <https://doi.org/10.1145/3663548.3675631>
- Accessibility Standards Canada. (2025). *CAN-ASC-6.2:2025, Accessible and equitable artificial intelligence systems*. <https://accessible.canada.ca/creating-accessibility-standards/asc-62-accessible-equitable-artificial-intelligence-systems>
- American Foundation for the Blind. (2025, January). *Guiding principles for more disability-inclusive AI*. <https://afb.org/research-and-initiatives/empowering-or-excluding/guiding-principles-more-disability-inclusive-ai>
- 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>
- Bickenbach, J. (2014). The capability approach and the International Classification of Functioning. *ALTER: European Journal of Disability Research*, 8, 10 to 23.
- Brand, D., Fischer Mogensen, K., & Engelbrecht, M. (2026, September 15). AI as reasonable accommodation and disability employment. *Disability & Society*. <https://doi.org/10.1080/09687599.2026.2726868>
- Burchardt, T. (2004). Capabilities and disability: The capabilities framework and the social model of disability. *Disability & Society*, 19(7), 735 to 751.
- Coffey, K. V., Krahn, G. L., Hanley, J. P., & Neely, J. E. (2026). Identifying implicit bias in LLM-based chat AI toward people with intellectual disabilities. *Disability and Health Journal*.
- CORDIS. (n.d.). *AIDE: Adaptive multimodal interfaces to assist disabled people in daily activities* (Horizon 2020 project 645322). European Commission. <https://cordis.europa.eu/project/id/645322>
- Della Penna, G., Buzzi, M., & Leporini, B. (2026). Generative AI as a new assistive technology for web interaction. In *Human-Computer Interaction, INTERACT 2025* (Lecture Notes in Computer Science, Vol. 16108). Springer. https://doi.org/10.1007/978-3-032-04999-5_6
- Davis Wright Tremaine LLP. (2026, September). *New technical standard for European Accessibility Act compliance released*. <https://www.dwt.com/insights/2026/09/european-accessibility-act-ict-standards-update>

- Demers, L., Weiss-Lambrou, R., & Ska, B. (2002). The Quebec User Evaluation of Satisfaction with Assistive Technology (QUEST 2.0). *Technology and Disability*, 14, 101 to 105.
- Deque Systems. (2026, September). *EN 301 549 v4.1.1 is final: What changed, what it means, and what you should do*. <https://www.deque.com/blog/en-301-549-v4-1-1-is-final-what-changed-what-it-means-and-what-you-should-do/>
- DeRosier, R., & Farber, R. (2005). Speech recognition software as an assistive device: A pilot study of user satisfaction and psychosocial impact. *Work*, 25(2), 125 to 134.
- Eberhardt, K. E. (2026, July 7). *Position paper: AI as a disability accommodation in higher education. Legal analysis and annotated bibliography on cognitive scaffolding and generative AI*. Zenodo. <https://doi.org/10.5281/zenodo.21231449>
- Ede, R. (2026, August 20). A thinking partner: Copilot as assistive technology. *Defra digital, data, technology and security blog*. <https://defradigital.blog.gov.uk/2026/08/20/a-thinking-partner-copilot-as-assistive-technology/>
- El Morr, C., Kundi, B., Mobeen, F., Taleghani, S., El-Lahib, Y., & Gorman, R. (2024). AI and disability: A systematic scoping review. *Health Informatics Journal*, 30(3).
- ETSI. (2026). *EN 301 549 V4.1.1 (2026-09), Accessibility requirements for ICT products and services*. https://www.etsi.org/deliver/etsi_en/301500_301599/301549/04.01.01_60/en_301549v040101p.pdf
- Gilly, T. (2026a, August). *Accommodation repricing: Disability, generative AI subscriptions, and the accessibility tax as a variable the vendor controls (v0.7)*. Real Safety AI Foundation. SSRN 7218538. <https://doi.org/10.2139/ssrn.7218538>
- Gilly, T. (2026b). *The accommodation axis: Disability discrimination, algorithmic opacity, and the deployers that anti-discrimination scholarship overlooks*. SSRN 7011798. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7011798
- Giraldo, M., & Sacchi, F. (2026). From technical performance to assistive validity: A critical review of evaluation approaches for AI-enabled assistive technology. *Healthcare*, 14(18), 2900. <https://doi.org/10.3390/healthcare14182900>
- Glazko, K., Yamagami, M., Desai, A., Mack, K. A., et al. (2023). *Three-month autoethnography of generative AI use for accessibility* (arXiv:2308.09924). <https://arxiv.org/abs/2308.09924>
- Gonzalez Penuela, R. E., Jung, C., Lin, S., Hu, R., & Azenkot, S. (2026). How multimodal large language models support access to visual information: A diary

study with blind and low vision people. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*.
<https://doi.org/10.1145/3772318.3793266>

- Hadar Souval, D., Haber, Y., Tal, A., Simon, T., Elyoseph, T., & Elyoseph, Z. (2025). Transforming perceptions: Exploring the multifaceted potential of generative AI for people with cognitive disabilities. *JMIR Neurotechnology*, 4, e64182. <https://doi.org/10.2196/64182>
- Hanuschock, W. E., & Silverman, A. M. (2026, August). *Innovation for access: AI and its use as assistive technology for people with disabilities*. American Foundation for the Blind. <https://afb.org/research-and-initiatives/ai-series/innovation-access>
- Jang, J., Moharana, S., & Carrington, P. (2024). “It’s the only thing I can trust”: Envisioning large language model use by autistic workers for communication assistance. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3613904.3642894>
- Johnson, J., Lewis, A., Mankoff, J., & Banner, O. (2026). “I don’t trust it, but I use it”: Navigating trust, privacy, and identity in disabled people’s use of generative AI. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3772318.3790652>
- Jutai, J., & Day, H. (2002). Psychosocial Impact of Assistive Devices Scale (PIADS). *Technology and Disability*, 14, 107 to 111.
- Knobel, N. (2026, April 10). Algorithmic risk and neuroinclusion: AI, psychosocial safety, and the rights of neurodivergent workers in Aotearoa New Zealand. *New Zealand Journal of Health and Safety Practice*, 3(1). <https://doi.org/10.26686/nzjhsp.v3i1.10606>
- Mitra, S. (2006). The capability approach and disability. *Journal of Disability Policy Studies*, 16(4), 236 to 247.
- Mladenov, T. (2025). AI and disabled people’s independent living: A framework for analysis. *AI & Society*. <https://doi.org/10.1007/s00146-025-02642-x>
- Moore, R. E., & Williams, A. B. (2020). AIDA: Using social scaffolding to assist workers with intellectual and developmental disabilities. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 366 to 368). <https://doi.org/10.1145/3371382.3378385>
- Moore, R. E., & Williams, A. B. (2021). Towards a learning architecture to support social scaffolding for an artificially intelligent disability assistant. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 467 to 469). <https://doi.org/10.1145/3434074.3447215>
- Morgan, D. (2026). Reasonable adjustments for neurodivergent students in higher education: Generative AI, accessibility, and mediated authorship. *Disability & Society*. <https://doi.org/10.1080/09687599.2026.2667528>

- National Disability Insurance Agency. (n.d.). *Framework for artificial intelligence-enabled assistive technology as supports under the NDIS*. <https://ndis.gov.au/news/8492-framework-artificial-intelligence-enabled-assistive-technology-supports-under-ndis>
- Nussbaum, M. C. (2006). *Frontiers of justice: Disability, nationality, species membership*. Harvard University Press. <https://doi.org/10.4159/9780674041578>
- 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-o>
- Randall, K. N., Drew, H., Gilman, E. S., & Dixon, E. (2025). Assistive technology uses and barriers in the home and workplace for adults with intellectual and developmental disabilities. *Journal of Applied Research in Intellectual Disabilities*, 38(1), e13306. <https://doi.org/10.1111/jar.13306> (first published online October 27, 2024)
- Phillips, B., & Zhao, H. (1993). Predictors of assistive technology abandonment. *Assistive Technology*, 5(1), 36 to 45. <https://doi.org/10.1080/10400435.1993.10132205>
- Polgar, J. M., Encarnação, P., Smith, E., & Cook, A. M. (2025). *Assistive technologies: Principles and practice* (6th ed.). Elsevier.
- Rodgers, R. (n.d.). AI can create plain-language and easy-read versions of every written material on earth. *Impact*, 38(3). Institute on Community Integration, University of Minnesota. <https://publications.ici.umn.edu/impact/38-3/ai-can-create-plain-language-and-easy-read-versions-of-every-written-material-on-earth-by-rylin-rodgers>
- Sustainable Bus. (2025, October). *Masats AIDA: Artificial intelligence bus doors*. <https://www.sustainable-bus.com/components/masats-aida-artificial-intelligence-bus-doors>
- Scherer, M. J., & Craddock, G. (2002). Matching Person and Technology (MPT) assessment process. *Technology and Disability*, 14(3), 125 to 131.
- Shew, A. (2023). *Against technoableism: Rethinking who needs improvement*. W. W. Norton.
- Shock, A., & Silverman, A. M. (2026, June). *Working with the machine: AI's expanding role in employment for people with disabilities*. American Foundation for the Blind. <https://afb.org/research-and-initiatives/ai-series/working-machine>
- Silverman, A. M., Whistler, A. L., Shock, A., Heydarian, C. H., Baguhn, S. J., Hanuschock, W. E., Hashimoto, M., Khan, O., & Vader, M.-L. (2026). *The AI quagmire: Benefits, risks, and user aspirations through a disability lens*. American Foundation for the Blind. <https://afb.org/AIResearch2>

- United Nations. (2006). *Convention on the Rights of Persons with Disabilities*. <https://www.un.org/disabilities/documents/convention/convoptprot-e.pdf>
- Silvera-Tawil, D., Higgins, L., Packer, K., Bayor, A. A., Walker, J. G., Li, J., Niven, P., Khanna, S., Byrnes, J., Bradford, D., et al. (2025). AI-enabled AT framework: A principles-based approach to emerging assistive technology. *Disability and Rehabilitation: Assistive Technology*, 20, 1955 to 1974. <https://doi.org/10.1080/17483107.2025.2479838>
- Valencia, S., Cave, R., Kallarackal, K., Seaver, K., Terry, M., & Kane, S. K. (2023). “The less I type, the better”: How AI language models can enhance or impede communication for AAC users. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3544548.3581560>
- van Toorn, G., Scully, J. L., & Gendera, S. (2025). “This robot is dictating her next steps in life”: Disability justice and relational AI ethics. *AI & Society*, 40. <https://doi.org/10.1007/s00146-025-02224-x>
- 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>
- Williams, K., & Breazeal, C. (2013). Reducing driver task load and promoting sociability through an Affective Intelligent Driving Agent (AIDA). In *Human-Computer Interaction, INTERACT 2013* (pp. 619 to 626). Springer.
- Williams, A. B., Williams, R. M., Moore, R. E., & McFarlane, M. (2019). AIDA: A social co-robot to uplift workers with intellectual and developmental disabilities. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction* (pp. 584 to 585). <https://doi.org/10.1109/HRI.2019.8673272>
- World Health Organization. (2001). *International classification of functioning, disability and health*. WHO.
- World Health Organization & UNICEF. (2022). *Global report on assistive technology*. WHO.
- World Wide Web Consortium. (2024). *Web Content Accessibility Guidelines (WCAG) 2.2*. <https://www.w3.org/TR/WCAG22/>
- Wright, C. (2026). Transcription is not generation: Distinguishing non-generative AI tool use from academic misconduct in higher education assessment. *International Journal for Educational Integrity*, 22, 24. <https://doi.org/10.1007/s40979-026-00234-w>
- Zhao, X., Cox, A., & Chen, X. (2025). The use of generative AI by students with disabilities in higher education. *The Internet and Higher Education*, 66, 101014. <https://doi.org/10.1016/j.iheduc.2025.101014>

- Zhao, X., Chen, X., & Cox, A. (2026). Exploring the affordances of generative AI in academic writing for students with disabilities. *Policy Futures in Education*, 24(1), 59 to 81. <https://doi.org/10.1177/14782103251395436>

Statutes and regulations

- 20 U.S.C. § 1401 (Individuals with Disabilities Education Act, definitions).
- 29 U.S.C. § 3002 (Assistive Technology Act, definitions).
- 42 U.S.C. § 12102 (Americans with Disabilities Act, definition of disability).
- 29 C.F.R. § 1630.2 (EEOC regulations implementing Title I of the ADA).
- Accessible Canada Act, S.C. 2019, c. 10.
- Disability Discrimination Act 1992 (Cth) (Australia).
- Equality Act 2010, c. 15 (United Kingdom).

Works by the author, with the sources supporting each

- **Puglisi, B. C. (n.d.-a). *HAIA: The Human Artificial Intelligence Assistant ecosystem* [Canonical framework page]. basilpuglisi.com. <https://basilpuglisi.com/haia-the-human-artificial-intelligence-assistant-ecosystem/>**
 - Atari, M., Xue, M. J., Park, P. S., Blasi, D. E., & Henrich, J. (2023). Which humans? PsyArXiv preprint. <https://doi.org/10.31234/osf.io/5b26t>
 - Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3), 61-83.
 - Puglisi, B. C. (2012). *Digital Facts: Twitter*. Digital Media Press.
 - Puglisi, B. C. (2025). *Governing AI: When Capability Exceeds Control*. ISBN 9798349677687.
 - Puglisi, B. C. (2025). *Ethics of artificial intelligence: A white paper on principles, risks, and responsibility*. Published August 18, 2025. <https://basilpuglisi.com/ethics-of-artificial-intelligence/>
 - Puglisi, B. C. (n.d.). *AI Provider Plurality: A congressional package* [AI Policy page]. <https://basilpuglisi.com/ai-provider-plurality-a-congressional-package/>
 - Puglisi, B. C. (n.d.). *Checkpoint-Based Governance* [Canonical framework page]. <https://basilpuglisi.com/cbg/>
 - Puglisi, B. C. (n.d.). *HAIA-CAIPR* [Canonical framework page]. <https://basilpuglisi.com/haia-caipr/>
 - Puglisi, B. C. (2026). *HAIA-RECCLIN case study 006 (v7)*. <https://basilpuglisi.com>
 - Puglisi, B. C. (n.d.). *HAIA-RECCLIN Reasoning and Dispatch* [Canonical framework page]. <https://basilpuglisi.com/haia-recclin/>

- Puglisi, B. C. (n.d.). *HEQ and AIS* [Canonical framework page]. <https://basilpuglisi.com/heq-ais/>
- Puglisi, B. C. (2026). *The loop that ate the governor* (v7). <https://basilpuglisi.com>
- Puglisi, B. C. (2026). *The minds that bend the machine: The voices shaping responsible AI governance* (forthcoming, October 2026). <https://basilpuglisi.com>
- **Puglisi, B. C. (n.d.-b). *HEQ and AIS: Measuring governance of AI literacy* [Canonical framework page]. basilpuglisi.com. <https://basilpuglisi.com/heq-ais/>**
 - Colorado General Assembly. (2026, May 14). *Senate Bill 26-189: Consumer protections in interactions with artificial intelligence systems*. Signed May 14, 2026; effective January 1, 2027.
 - European Parliament and Council. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act)*. Official Journal of the European Union.
 - European Parliament and Council. (2026). *Regulation (EU) 2026/1744 amending Regulation (EU) 2024/1689 (Digital Omnibus on AI)*. Official Journal of the European Union, July 24, 2026; in force July 27, 2026.
 - International Organization for Standardization. (2023). *ISO/IEC 42001:2023, Information technology, Artificial intelligence, Management system*. ISO.
 - Landgericht München I. (2026, May 28). *Case No. 26 O 869/26*. Preliminary injunction in expedited proceedings; on appeal.
 - Lokken v. UnitedHealth Group, Inc., No. 0:23-cv-03514 (D. Minn.). Discovery order, March 9, 2026.
 - National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)*. NIST AI 100-1.
 - Organisation for Economic Co-operation and Development. (2019, updated 2024). *OECD AI Principles*. oecd.ai
 - AERA, APA, & NCME. (2014). *Standards for educational and psychological testing*. American Educational Research Association.
 - Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>
 - Buçinca, Z., Malaya, M. B., & Gajos, K. Z. (2021). To trust or to think: Cognitive forcing functions can reduce overreliance on AI. *Proceedings of*

the ACM on Human-Computer Interaction, 5(CSCW1), Article 188.
<https://doi.org/10.1145/3449287>

- Ganuthula, V. R. R., & Balaraman, K. K. (2025). Artificial intelligence quotient framework for measuring human collaboration with artificial intelligence. *Discover Artificial Intelligence*, 5, Article 268.
<https://doi.org/10.1007/s44163-025-00516-1>
- Gerlich, M. (2025a). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), Article 6.
<https://doi.org/10.3390/soc15010006>
- Hollnagel, E., & Woods, D. D. (2005). *Joint cognitive systems: Foundations of cognitive systems engineering*. CRC Press.
- Hutchins, E. (1995). *Cognition in the wild*. MIT Press.
- Insurance Business. (2026, February 11). Chaucer, Armilla bet big on standalone AI coverage as reinsurers struggle to keep pace. [insurancebusinessmag.com](https://www.insurancebusinessmag.com)
- The Insurer. (2026, April 24). Standalone AI liability market takes shape with underwriting discipline key to MGA success. [theinsurer.com](https://www.theinsurer.com)
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50 to 80.
https://doi.org/10.1518/hfes.46.1.50_30392
- Legg, S., & Hutter, M. (2007). Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4), 391 to 444.
<https://doi.org/10.1007/s11023-007-9079-x>
- Lior, A. (2025). E/Insuring the AI age: Empirical insights into artificial intelligence liability policies. *Connecticut Insurance Law Journal*, 31, 99. SSRN Abstract 5316376.
- Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1 to 16). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376727>
- National Association of Insurance Commissioners. (2023, December 4). *Model bulletin on the use of artificial intelligence systems by insurers*. content.naic.org
- Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, Article 100041.
<https://doi.org/10.1016/j.caeai.2021.100041>
- OECD & European Commission. (2026). *Empowering learners for the age of AI: An AI literacy framework for primary and secondary education*. OECD Publishing. <https://doi.org/10.1787/65cd27d4-en>

- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676 to 688. <https://doi.org/10.1016/j.tics.2016.07.002>
- ScienceSoft. (2026, September 10). AI risks to enter 60 to 80% of liability and cyber insurance underwriting by 2028. GlobeNewswire.
- Sidra, S., & Mason, C. (2025). Generative AI in human-AI collaboration: Validation of the Collaborative AI Literacy and Collaborative AI Metacognition Scales for effective use. *International Journal of Human-Computer Interaction*. <https://doi.org/10.1080/10447318.2025.2543997>
- Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory. *Science*, 333(6043), 776 to 778. <https://doi.org/10.1126/science.1207745>
- U.S. Department of Labor. (2026). *Training and Employment Notice No. 07-25: The U.S. Department of Labor's Artificial Intelligence Literacy Framework*. <https://www.dol.gov/agencies/eta/advisories/ten-07-25>
- UNESCO. (2024). AI competency framework for students. United Nations Educational, Scientific and Cultural Organization. <https://doi.org/10.54675/JKJB9835>
- Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12), 2293 to 2303. <https://doi.org/10.1038/s41562-024-02024-1>
- Vasconcelos, H., Jörke, M., Grunde-McLaughlin, M., Gerstenberg, T., Bernstein, M. S., & Krishna, R. (2023). Explanations can reduce overreliance on AI systems during decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1), Article 129. <https://doi.org/10.1145/3579605>
- **Puglisi, B. C. (n.d.-c). *What is Checkpoint-Based Governance? Human authority [Canonical framework page]. basilpuglisi.com. https://basilpuglisi.com/cbg/***
 - Asimov, I. (1942). Runaround. *Astounding Science Fiction*.
 - Asimov, I. (1985). *Robots and Empire*. Doubleday.
 - Puglisi, B. C. (2026). HAIA-CARCS: Compliance accountability record and case study. basilpuglisi.com.
 - Puglisi, B. C. (2026). Why you cannot program or prompt governance into AI. basilpuglisi.com.
 - Puglisi, B. C. (2026). HAIA specifications. <https://github.com/basilpuglisi>
- **Puglisi, B. C. (2026a, January 1). *Ethical AI, Responsible AI, and AI Governance are not the same thing. basilpuglisi.com. https://basilpuglisi.com/what-we-failed-to-define-is-how-we-fail/***

- Bai, Y., et al. (2022). Constitutional AI: Harmlessness from AI feedback. arXiv. <https://arxiv.org/abs/2212.08073>
- Center for AI Safety. (2023, May 30). Statement on AI risk. <https://safe.ai/work/press-release-ai-risk>
- CGTN. (2024, December 28). 30 years left? AI ‘Godfather’ warns the technology may end humanity. <https://newseu.cgtn.com/news/2024-12-28/AI-Godfather-warns-rapid-development-can-cause-human-extinction-1zHEIq63EDC/index.html>
- Dietterich, T. G. (2000). Ensemble methods in machine learning. In Multiple classifier systems (pp. 1-15). Springer. <https://web.engr.oregonstate.edu/~tgd/publications/mcs-ensembles.pdf>
- European Commission. (2019). Ethics guidelines for trustworthy AI. High-Level Expert Group on Artificial Intelligence. https://www.europarl.europa.eu/cmsdata/196377/AI%20HLEG_Ethics%20Guidelines%20for%20Trustworthy%20AI.pdf
- European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act), Article 14: Human oversight. Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- Harper, D. (n.d.). Govern. In Online Etymology Dictionary. <https://www.etymonline.com/word/govern>
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3), 61-83. <https://www2.psych.ubc.ca/~henrich/pdfs/WeirdPeople.pdf>
- IBM. (2025). AI governance. IBM Think. <https://www.ibm.com/think/topics/ai-governance>
- International Organization for Standardization. (2021). ISO 37000:2021 Governance of organizations. <https://www.iso.org/standard/65036.html>
- International Organization for Standardization. (2023). ISO/IEC 42001:2023 AI management systems. <https://www.iso.org/standard/42001>
- Irving, G., Christiano, P., & Amodei, D. (2018). AI safety via debate. arXiv. <https://arxiv.org/abs/1805.00899>
- Tao, Y., Viberg, O., Baker, R. S., & Kizilcec, R. F. (2024). Cultural bias and cultural alignment of large language models. *PNAS Nexus*, 3(9), pga346. <https://pmc.ncbi.nlm.nih.gov/articles/PMC11407280/>
- Merriam-Webster. (n.d.). Govern. In Merriam-Webster.com dictionary. <https://www.merriam-webster.com/dictionary/govern>
- Merriam-Webster. (n.d.). Governance. In Merriam-Webster.com dictionary. <https://www.merriam-webster.com/dictionary/governance>

- Microsoft. (2022). Microsoft Responsible AI Standard v2: General requirements. <https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/final/en-us/microsoft-brand/documents/Microsoft-Responsible-AI-Standard-General-Requirements.pdf>
- MIT Sloan School of Management. (2023, May 23). Why neural net pioneer Geoffrey Hinton is sounding the alarm on AI. <https://mitsloan.mit.edu/ideas-made-to-matter/why-neural-net-pioneer-geoffrey-hinton-sounding-alarm-ai>
- National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). U.S. Department of Commerce. <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>
- National Institute of Standards and Technology. (2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1). U.S. Department of Commerce. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- OECD. (2019). Recommendation of the Council on Artificial Intelligence. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- OpenAI. (2023, July 5). Introducing Superalignment. <https://openai.com/index/introducing-superalignment/>
- Oxford English Dictionary. (n.d.). Govern, v. In OED Online. Oxford University Press. https://www.oed.com/dictionary/govern_v
- PBS. (2023, May 9). Geoffrey Hinton warns of the “existential threat” of AI. Amanpour and Company. <https://www.pbs.org/video/godfather-of-ai-warns-of-the-existential-threat-of-ai-lj1i1c/>
- UNESCO. (2021). Recommendation on the Ethics of Artificial Intelligence. <https://unesdoc.unesco.org/ark:/48223/pf0000380455>
- **Puglisi, B. C. (2026b, May). *The other AI: Augmented intelligence and the honest future of human-AI collaboration*. basilpuglisi.com.**
 - Anthropic. (2025a, July). System card: Claude Opus 4 & Claude Sonnet 4. <https://www-cdn.anthropic.com/4263b940cabb546aa0e3283f35b686f4f3b2ff47.pdf>
 - Anthropic. (2025b, June). Agentic misalignment: How LLMs could be insider threats. <https://www.anthropic.com/research/agentic-misalignment>
 - Atari, M., Xue, M. J., Park, P. S., Blasi, D. E., & Henrich, J. (2023). Which humans? PsyArXiv preprint. <https://doi.org/10.31234/osf.io/5b26t>
 - Bartlett, S. (Host). (2025, November 27). AI expert: Here is what the world looks like in 2 years! Tristan Harris [Video]. *The Diary of a CEO*. <https://www.youtube.com/watch?v=BFU1OckhBwo>

- Bartlett, S. (Host). (2026, March 26). AI whistleblower: We are being gaslit by the AI companies! Karen Hao [Video]. *The Diary of a CEO*. <https://www.youtube.com/watch?v=Cn8HBj8QAbk>
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610 to 623. <https://doi.org/10.1145/3442188.3445922>
- Brynjolfsson, E., Chandar, B., & Chen, R. (2025). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence [Working paper]. Stanford Digital Economy Lab. <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>
- Chalk and Talk Podcast. (2026, January 9). Why more classroom technology is making students learn less [Audio interview with John Horvath, hosted by Anna Stokke]. *Chalk and Talk*. <https://www.stokkemath.com/chalk-and-talk-podcast/>
- Dellermann, D., Ebel, P., Sollner, M., & Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, 61(5), 637 to 643.
- Dweck, C. S. (2006). *Mindset: The new psychology of success*. Random House.
- Engelbart, D. C. (1962). *Augmenting human intellect: A conceptual framework*. Stanford Research Institute.
- EY. (2025, June 4). *EY Responsible AI Pulse Survey: AI adoption outpaces governance as risk awareness among the C-suite remains low*. https://www.ey.com/en_gl/newsroom/2025/06/ey-survey-ai-adoption-outpaces-governance-as-risk-awareness-among-the-c-suite-remains-low
- Hao, K. (2025). *Empire of AI: Dreams and nightmares in Sam Altman's OpenAI*. Penguin Press.
- Harvard Crimson. (2023, May 22). More than three-quarters of surveyed Harvard faculty identify as liberal. *The Harvard Crimson*. <https://www.thecrimson.com/article/2023/5/22/faculty-survey-2023-politics/>
- Hemmer, P., Schemmer, M., Köhl, N., Vössing, M., & Satzger, G. (2025). Complementarity in human-AI collaboration: Concept, sources, and

- evidence. *European Journal of Information Systems*, 34(6), 979 to 1002.
<https://doi.org/10.1080/0960085X.2025.2475962>
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The WEIRD people in the world? *Behavioral and Brain Sciences*, 33(2-3), 61 to 83.
<https://doi.org/10.1017/S0140525X0999152X>
 - Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
 - Kasparov, G. (2017). *Deep thinking: Where machine intelligence ends and human creativity begins*. PublicAffairs.
 - Lane, J. N., Boussioux, L., Ayoubi, C., Chen, Y. H., Lin, C., Spens, R., Wagh, P., & Wang, P. H. (2025). The narrative AI advantage? A field experiment on AI-augmented evaluations of early-stage innovations. Harvard Business School Working Paper No. 25-001. (Revised May 2025.)
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4914367
 - Licklider, J. C. R. (1960). Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*, HFE-1, 4 to 11.
<https://doi.org/10.1109/THFE2.1960.4503259>
 - Marcus, G., & Davis, E. (2019). *Rebooting AI: Building artificial intelligence we can trust*. Pantheon Books.
 - McKinsey & Company. (2025, March 11). The state of AI: Global survey.
<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
 - Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187 to 192.
<https://doi.org/10.1126/science.adh2586>
 - Puglisi, B. C. (2012). *Digital Factics: Twitter* (Factics methodology origin). MagCloud. basilpuglisi.com.
 - Puglisi, B. C. (2025). The Growth OS: Parts 1 and 2. basilpuglisi.com.
 - Puglisi, B. C. (2025). *Governing AI: When capability exceeds control* (ISBN 9798349677687). Chapter 2: Corporate Incentives and Economics. basilpuglisi.com.
 - Puglisi, B. C. (n.d.). *Checkpoint-Based Governance* [Canonical framework page]. <https://basilpuglisi.com/cbg/>
 - Puglisi, B. C. (2026). Empire of evidence: Testing Karen Hao's claims against the governance infrastructure they require. basilpuglisi.com.
<https://basilpuglisi.com/empire-of-evidence-testing-karen-hao-claims-governance-infrastructure/>
 - Puglisi, B. C. (2026). AI governance beyond the warning: From Tristan Harris's diagnosis to the infrastructure it requires. basilpuglisi.com.
<https://basilpuglisi.com/ai-governance-beyond-the-warning/>

- Puglisi, B. C. (2026). *AI as a mirror to humanity: Do what we say, not what we do*. basilpuglisi.com.
- Puglisi, B. C. (2026). Bridging the measurement gap in augmented intelligence: The HEQ and AIS. SSRN Abstract 6583419. <https://ssrn.com/abstract=6583419>
- Puglisi, B. C. (n.d.). *HAIA-RECCLIN Reasoning and Dispatch* [Canonical framework page]. <https://basilpuglisi.com/haia-recclin/>
- Puglisi, B. C. (2026). HAIA-RECCLIN Case Study 008: HEQ longitudinal administration and rubric discovery. basilpuglisi.com.
- Puglisi, B. C. (2026). The AI cognitive decline narrative has not tested what it claims: A methodological audit. basilpuglisi.com. <https://basilpuglisi.com/ai-cognitive-decline-narrative-untested/>
- Puglisi, B. C. (2026). How credentialed professionals shape policy when method governance is stripped: The Horvath case study. basilpuglisi.com. <https://basilpuglisi.com/how-credentialed-testimony-outpaces-research-horvath-case-study/>
- Puglisi, B. C. (2026). Human drift and hallucination: The data literacy crisis hiding behind the AI one. basilpuglisi.com. <https://basilpuglisi.com/human-drift-and-hallucination-the-data-literacy-crisis-hiding-behind-the-ai-one/>
- Puglisi, B. C. (n.d.). *AI Provider Plurality: A congressional package* [AI Policy page]. <https://basilpuglisi.com/ai-provider-plurality-a-congressional-package/>
- Puglisi, B. C. (n.d.). *HAIA-CAIPR* [Canonical framework page]. <https://basilpuglisi.com/haia-caipr/>
- PwC. (2025). *Global AI Jobs Barometer*. <https://www.pwc.com/gx/en/issues/artificial-intelligence/ai-jobs-barometer.html>
- Resume.org. (2025, September). *AI and workforce displacement survey*. <https://www.resume.org>
- Tetlock, P. E., & Gardner, D. (2015). *Superforecasting: The art and science of prediction*. Crown Publishers.
- Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12), 2293 to 2303. <https://doi.org/10.1038/s41562-024-02024-1>
- World Economic Forum. (2025). *Future of Jobs Report 2025*. <https://www.weforum.org/publications/the-future-of-jobs-report-2025>

Disclaimer

I am not a lawyer, and this article does not provide legal advice. This is thought research and governance analysis based on public sources, cited materials, and human-AI review. It is intended to help executives, practitioners, insurers, and governance teams think more clearly about AI risk, liability exposure, and documentation practices. Readers should not rely on this article as a legal opinion, compliance determination, or substitute for qualified counsel. Any organization facing a legal, regulatory, contractual, or insurance question should consult its own attorney, broker, or professional adviser before acting.

The Other AI: Audio Briefings on Augmented Intelligence and AI Governance

[Spotify](#) | [Apple Podcasts](#) | [Amazon Music](#) | [YouTube Playlist](#)

#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.