IIAR JOURNALS

Innovating Knowledge. Advancing Research.

IIAR Journal of Disability, Inclusion, and Assistive Technologies

Vol 1, No 1, 2026

Artificial Intelligence as Assistive Technology: A Theoretical Framework of Capability Extension, Dependence and Dignity for Persons with Disabilities

IGWE Emmanuel Sunday[1], NWOGBOJI Emmanuel Usulor[2], BAMIGBOLA Paul Idowu[3]

[1,3] Veristzon Global Research Lab; [2] International Institute of AI and Research

ABSTRACT

Artificial intelligence has turned ordinary consumer devices into assistive technologies. A phone now describes a street scene to a blind user, a cloud service captions a conversation for a deaf colleague, a decoder guesses the intended movement of a paralysed hand, and a chatbot helps an adult with intellectual disability think through a decision. The research literature evaluates these systems almost entirely on functional performance, and the problem this paper addresses is that functional performance says little about whether a disabled person's life has become larger, more secure or more their own. The purpose of the paper is to develop a framework for evaluating AI-based assistive technology beyond accuracy. It is a theoretical study. Forty peer-reviewed journal articles were retrieved under a 2023–2026 search window, primarily through PubMed and supplemented by web search; all forty retained articles date from 2024 to 2026. Each was verified against publisher records, screened for retractions, and read against the capability approach, the social model of disability and the human–AI interaction literature. The analysis distinguishes three effects any assistive AI produces at once (extending capability, creating dependence and shaping dignity) and shows that they routinely move in opposite directions: the same system that lets a person do more may bind them to a vendor, a data connection or a subscription, and may do so in a register they experience as diminishing. Eight propositions specify the design and provisioning conditions under which capability extension coincides with autonomy and dignity. A further analysis identifies the risks that AI adds to the older problems of assistive technology provision: probabilistic error in safety-critical settings, surveillance, obsolescence and affordability. The framework is situated in African provisioning conditions, where the distance between what the technology could do and what reaches disabled people is widest. Recommendations are directed at ministries, regulators, procurers, universities, designers and organisations of persons with disabilities.

Keywords: assistive technology; artificial intelligence; disability; capability approach; dependence

Artificial Intelligence as Assistive Technology: A Theoretical Framework of Capability Extension, Dependence and Dignity for Persons with Disabilities

1. INTRODUCTION

Assistive technology used to be recognisable. A white cane, a hearing aid, a wheelchair, a Perkins brailler: each was a discrete object, built for a defined impairment, supplied through a defined channel and paid for, when it was paid for at all, by a health or welfare budget. Artificial intelligence has dissolved that recognisability. The most consequential assistive technologies of the last three years are features inside general-purpose products: a vision-language model that answers a blind person's questions about a photograph, a speech recogniser that produces live captions, a predictive interface that completes what a person with a motor impairment has begun to type. None of these was procured by a rehabilitation service, none is listed on a national priority assistive products list, and most of them change behaviour whenever the vendor updates the model.

That shift has produced a great deal of engineering and surprisingly little theory. Scoping reviews of AI for disability now cover more than two hundred studies between them (Alsaleh, 2025; Du et al., 2026; Kaboré et al., 2026), and they converge on an uncomfortable picture: most studies originate in high-income countries (Kaboré et al., 2026), the field favours functional over psychosocial support and offers limited real-world evidence of benefit (Du et al., 2026), and access remains unequal, particularly in low-resource settings (Alsaleh, 2025). Kaboré et al. (2026) found that of 93 health-related AI studies concerning disabled people, only six actively involved them through co-design or user validation. Giraldo and Sacchi (2026), reviewing sixteen evaluation approaches for AI-enabled assistive technology, found that only three explicitly evaluated the properties specific to AI (its probabilistic behaviour, its dependence on data, its tendency to adapt) and that none combined such evaluation with comprehensive coverage of the person, their activities, their participation and their environment. The gap this paper addresses is not a lack of evidence that AI assistive technology works. It is the absence of a way of saying what "works" should mean.

The paper's purpose is to supply that vocabulary. It argues that any AI-based assistive technology produces three effects simultaneously, and that evaluation must hold all three in view. The first is capability extension: the system enlarges the set of things a person can actually do and be. The second is dependence: the system introduces new relations of reliance on vendors, infrastructures, data flows and update cycles that the person did not previously have and cannot easily exit. The third is dignity: the system shapes how the person is treated and how they experience themselves while being assisted, in ways that may be enabling or demeaning independently of what the system achieves. The central claim is that these effects are separable and often move in opposite directions. A system can extend capability while deepening dependence and corroding dignity, and a weaker system can do less while leaving the person more fully in command of their life.

The argument matters most where provision is weakest. In sub-Saharan Africa fewer than three per cent of those who need an assistive product have one (Pakouyowou, 2026), and the products that do arrive are costly, largely imported and constrained by weak local production (Gaiotto et al., 2026). AI-based assistance promises to bypass that bottleneck because it rides on smartphones and cloud services that are already present. But the same features that make it reach further also make its dependence and dignity effects sharper. Africa is not an afterthought for this framework but its most demanding test.

Three questions organise the work. Under what conditions does AI-based assistive technology extend what a disabled person can do? When does that extension generate dependence that the person cannot control? And what determines whether being assisted by an AI is experienced as dignified?

2. LITERATURE REVIEW

2.1 CONCEPTUAL REVIEW

Three concepts carry the argument and each needs to be defined with more care than the engineering literature gives it.

The first is assistive technology itself. The World Health Organization's umbrella definition covers products and the services needed to deliver them, and national assessments now use its "5P" framework of people, providers, personnel, policy and products (Joskow et al., 2025). On that view a product without a service system is not assistive technology in any meaningful sense. AI-based assistance strains this

definition at both ends. The assistive function is usually a feature of something else, such as an operating system, a messaging application or a wearable, so there is no device to list, procure or repair, and provision is handled by app stores and subscription billing rather than by assessors and fitters. Recent work on digital assistive technology in low- and middle-income countries treats the smartphone as one of several delivery platforms rather than the default (Wade et al., 2026), and this paper follows that lead: what matters is not whether a product counts as assistive technology but which functions it performs for whom, under what terms.

The second concept is capability. Following the capability approach as it has been applied to disability (Joskow et al., 2025; Ortiz-Escobar et al., 2026; Ryan & O'Brien, 2024), a capability is a real opportunity to do or be something a person has reason to value, as distinct from a functioning, which is the doing or being itself. The distinction matters for assistive technology because a device can expand functionings without expanding capabilities: a person who can now read a document only when the connection is up and the subscription is paid has gained a functioning that is contingent on conditions outside their control. Ortiz-Escobar et al. (2026) make the point sharply using the concept of "insecure functioning" from disadvantage theory: in their study of visually disabled adults in Colombia, valued functionings such as employment and mobility often remained insecure because of systemic barriers, and a functioning that could be withdrawn at any moment is worth less than the same functioning held securely. On this paper's reading, much of what AI-based assistance delivers is insecure in exactly this way.

The third concept is dignity. The term is used in disability rights instruments without definition, and the paper does not attempt a philosophical account. It uses dignity in a practical sense: the degree to which an arrangement treats a person as an agent whose purposes matter, rather than as a problem to be managed, and the degree to which the person experiences it that way. Two recent findings illustrate the practical meaning. Coffey et al. (2026) prompted five large language models to generate 25,000 stories with and without descriptors for intellectual disability and found that characters with intellectual disability were written as younger, more dependent, more in need of rescue and more "inspirational". A conversational agent that carries those priors into an assistive role will be condescending in ways no accuracy metric records. By contrast, Álvarez-Aguado et al. (2026) found that adults with intellectual disability in Chile valued a conversational AI precisely because it helped them organise their thinking while leaving the decision with them. Dignity, in this usage, is a property of the relation between the person and the assistance, not of the technology alone.

Dependence needs a brief note. Everyone depends on infrastructure, and disability studies has long argued that independence is a misleading ideal. The relevant question is whether dependence is visible, exitable, reversible and distributed rather than concentrated in a single actor. Dependence on a cane is visible and exitable, and withdrawing the cane leaves the person no worse off than before it; dependence on a cloud model that may be withdrawn, repriced or silently changed is none of these.

2.2 THEORETICAL REVIEW

The framework draws on three bodies of theory.

The capability approach supplies the evaluative standard. Its application to assistive technology has moved in the last two years from philosophical invocation to empirical operationalisation. Joskow et al. (2025) applied it to England's 2023 national assessment of assistive technology, drawing on a survey of 7,000 disabled adults and children, and identified three domains of impact that users themselves report: functions and activities, outcomes such as autonomy and safety, and lived experience including a sense of control and "peace of mind". Ryan and O'Brien (2024) used the approach with 50 people with intellectual disabilities and 28 family carers and concluded that people with intellectual disabilities can lead "capability-deprived" lives, with "lowered expectations and static lives" among the themes their participants described. The approach is useful here because it separates means from ends: a model's accuracy is a means, and the question is always what real opportunities it opens.

The social model of disability and the disability rights framework codified in the Convention on the Rights of Persons with Disabilities supply the critical standard. Their contribution is the insistence that disability is produced by the mismatch between bodies and environments rather than by impairment, and that remedies which fix the person while leaving the environment untouched are suspect. Bariffi (2026) applies this to neurotechnology, arguing that brain-computer interfaces and cognitive monitoring tools, deployed uncritically, risk replicating the coercion and paternalism of the medical model under a new name, and that governance must move "from enhancement and correction to inclusion and

empowerment". Dalal et al. (2026) make a parallel argument about AI for disability health equity: most models rely on biomedical and functional data and are trained on datasets that neglect the social and structural determinants of disability, and so insufficiently capture capabilities, recognition or structural justice. Mahmoudi (2026) addresses the same tendency under the heading of techno-solutionism, from a biopsychosocialtech perspective on assistive technology. For this paper, the social model supplies a warning that capability extension through AI can be a way of making the individual bear the cost of inaccessible environments.

The human–AI interaction literature supplies the mechanism. Its most relevant concepts are shared autonomy, calibrated trust and reliance. Lee et al. (2025) demonstrate shared autonomy directly: an AI "copilot" collaborating with a non-invasive brain-computer interface allowed a participant with paralysis to achieve a 3.9-times higher target hit rate and to complete a robotic pick-and-place task that was impossible without the copilot. The same study makes the theoretical problem vivid, because the copilot infers the user's goal and acts on the inference, so the line between assisting the person's intention and substituting for it is a design parameter rather than a fact. Alkhateeb (2026), following 547 adults with dyslexia over four weeks, found that anxiety reduced "adaptive reliance" on generative AI while improvements in metacognitive regulation increased it, which suggests that whether reliance is healthy depends on the person's capacity to monitor it. These findings give the framework its hinge: dependence is not an inherent property of an AI system but an outcome of how control, inference and feedback are distributed between the person and the machine.

2.3 EMPIRICAL REVIEW

The empirical literature can be read as a series of demonstrations that capability extension is real, followed by a series of demonstrations that the extension is conditional in ways the first set rarely acknowledges.

For blind and low-vision users, vision-language models are the clearest case of extension. Arystanbekov et al. (2026) built a wearable system running a quantised LLaVA-NeXT model entirely on-device, validated it against the VizWiz benchmark of photographs taken by blind people, and reported a 25 per cent increase in image identification accuracy over a captioning baseline with 28 participants. Singh et al. (2026) argue that general-purpose vision-language models are poor at conveying spatial information and built smart glasses that use open-vocabulary object detection to guide a blind user's hand to an object. Mekonnen et al. (2025) showed that a Raspberry Pi system combining object detection, optical character recognition and offline voice control reached 75 to 90 per cent accuracy without any cloud connection. Arystanbekov et al. (2026) position their system against cloud-based captioning that suffers latency and dependence on connectivity, and Mekonnen et al. (2025) against cloud processing that raises privacy concerns and limits access in resource-constrained settings, which is to say that the dominant commercial pattern of AI assistance for blind users is one the engineers themselves regard as fragile.

For deaf, hard-of-hearing and speech-impaired users, the evidence is more sobering. Zhao et al. (2025) tested four commercial speech-to-text services on 850 recordings and found a mean word error rate of 52.6 per cent for deaf and hard-of-hearing speakers against 5.0 per cent for hearing speakers; for speakers with prelingual hearing loss the rate was 80.5 per cent, and for those who primarily communicate in sign language it was 70.2 per cent. Urqueta Alfaro et al. (2026) found that caption accuracy for accented speech was worse on both word error and semantic similarity measures across three engines. Personalised models help but do not close the gap: Tobin et al. (2024) found a persistent accuracy loss between read and conversational speech even in personalised models for 27 speakers with speech disorders, and Cave (2024), following three people with amyotrophic lateral sclerosis using Google's Project Relate for up to a year, reported that all of them found accuracy in real conversation lower than expected and useful only in certain contexts. The pattern is instructive: the users who most need speech recognition are the ones it serves worst, and personalisation ties the user's communication to a single vendor's model.

For motor impairment, the brain-computer interface literature shows extension at its most dramatic and dependence at its most literal. Oxley et al. (2025), reflecting on the 2023 International BCI Society meeting, identify four requirements for user autonomy (immediate use, ease of use, continuous use and stable system use), and the last two are precisely the properties that depend on a commercial entity continuing to exist and to support the device. Ashouri et al. (2026) note "recent cases of consumer exploitation, misleading claims, and inadequate patient aftercare" in neurotechnology, which is the obsolescence problem stated in its starkest form: a

person whose capability depends on an implant is dependent on a company.

For cognitive and intellectual disability, the evidence is thin but pointed. Álvarez-Aguado et al. (2026) found that the legitimacy of a conversational AI depended on its being "accessible, personalised, and user-controlled" and that participants rejected substitution for human support, which is the best available statement of what dignified assistance looks like from the user's side. Marul (2026), mapping aided augmentative and alternative communication for people with multiple disabilities, found the evidence for AI functions in this population "critically limited".

Turning to provision, the African literature describes the conditions under which any of this would have to work. Acheampong et al. (2026) surveyed 189 people with mobility and visual impairments in urban Ghana and found that about half of those with mobility impairments and 44 per cent of those with visual impairments identified high device costs as a major barrier, and that women used assistive technology less than men (43.4 against 56.6 per cent). Zandam et al. (2025), analysing Nigeria's 2021 Multiple Indicator Cluster Survey of 22,443 children, found that only 43.2 per cent of children with seeing difficulties, 37.2 per cent with hearing difficulties and 28.9 per cent with walking difficulties used an assistive device. Kibret et al. (2025) describe students in Ethiopia facing inaccessible facilities, a lack of learning materials in accessible formats and negative attitudes, Kamenopoulou et al. (2026) describe deafblind adults in Kenya relying significantly on human support alongside the assistive products they used, and Ogondiek (2026) documents the gap between Tanzania's policy commitments and what reaches pupils. Hendy et al. (2026), interviewing women with physical and visual disabilities in Minya, Egypt, found that AI-enhanced glasses, smart canes and screen readers were perceived as enabling but that high costs, limited maintenance services, privacy concerns and "economic dependency" stood in the way. At the system level, Anderson et al. (2026a, 2026b) and Curati-Alasonatti et al. (2026) report on national coordination mechanisms, priority product lists and ecosystem programmes across sub-Saharan Africa and other low- and middle-income countries, while Wilkens et al. (2026) and Gaiotto et al. (2026) analyse financing and tariff options. Apart from the Egyptian interviews, none of this provisioning literature yet has a category for AI-based assistance delivered through consumer platforms, and the AI literature barely engages with provisioning at all. The two bodies of work describe the same people and seldom cite each other.

Three further findings bear on dignity. Kim and Park (2026), using South Korea's 2023 Digital Divide Survey with 2,200 disabled respondents, found that a subgroup perceived themselves as more exposed to AI risks than society at large, a pattern they call self-relative pessimism. Soldatic et al. (2026) found that care providers supporting migrants with disabilities experienced AI as emotionally charged, fearing the loss of relational care. Umbelino et al. (2026), reviewing 95 studies of assistive technology design, found that users were seldom involved before the testing phase, by which point the decisions that determine dignity have already been made.

3. MATERIALS AND METHODS

This is a theoretical paper and its method is conceptual synthesis: the construction of a framework from existing concepts and empirical findings, followed by the derivation of propositions that could be tested. No human participants were involved, no primary data were collected, and ethical approval was not required. The only dataset is the reference list, which is openly available through the digital object identifiers given for each source.

The literature was assembled in October 2026. The primary retrieval tool was PubMed, searched with title- and abstract-restricted Boolean queries combining terms for artificial intelligence, large language models, speech recognition, vision-language models and brain-computer interfaces with terms for disability, assistive technology, the capability approach, privacy and the names of African countries. Web search was used to look for candidate articles outside the biomedical index, although most of what it returned were preprints and conference proceedings and were excluded. Inclusion criteria were peer-reviewed journal articles published between 2023 and 2026 in English with a registered digital object identifier. Preprints, conference papers, book chapters, dissertations and reports were excluded without exception; the Convention on the Rights of Persons with Disabilities is referred to only through a 2026 legal analysis of it. Approximately 180 records were screened at title and abstract level across eleven searches; 40 were retained after reading the abstract and, where open access permitted, the full text.

Each retained record was verified against its publisher metadata as held in PubMed and Crossref: author names, title, journal, year, volume, issue and

pagination were copied from those records rather than from memory or search snippets. Retraction screening inspected the Crossref record for update-to and relation fields and the PubMed record for the "Retracted Publication" type. This mattered in practice. One search returned a 2025 special issue of a leading assistive technology journal, titled "AI empowered Assistive Technology", every article of which had been retracted by the publisher for compromised peer review; those articles and two other retractions encountered were discarded. Any reader can repeat the screen using the digital object identifiers in the reference list.

The analytic procedure had three steps. Each source was coded for which of three effects it documented (capability extension, dependence or dignity) and for the mechanism through which the effect arose. The mechanisms were then grouped into design conditions (properties of the system) and provisioning conditions (properties of the channel through which it reaches the person). Propositions were written wherever two or more independent sources, from different populations or disciplines, pointed to the same conditional relationship, and were phrased so that a future empirical study could disconfirm them. The basis of judgement throughout was convergence across sources rather than the strength of any single one; where the evidence is thin, as it is for cognitive disability and for most of Africa, the propositions are flagged as such in Section 4.2.

An artificial intelligence writing assistant was used, under the author's direction, to support literature retrieval, metadata verification, drafting and formatting. Every source cited was retrieved and checked against its publisher record as described above, and the author is responsible for the selection of sources, the argument and the text.

4. RESULTS AND DISCUSSION

4.1 RESULTS

The framework is built on a single observation: an AI-based assistive technology is never only a capability-extending device. It is also a set of dependencies and a social relation. Table 1 sets out the three effects, their definitions, the dimensions along which each varies and the sources that ground them.

Table. 1. The capability–dependence–dignity framework for AI-based assistive technology.

Effect	Definition	Dimensions along which it varies	Illustrative sources
Capability extension	Enlargement of the set of real opportunities to do and be what the person values, as distinct from task performance on a benchmark	Breadth (how many valued activities are opened); security (whether the opportunity holds when conditions change); conversion (whether the person's circumstances allow the function to become an opportunity)	Joskow et al. (2025); Ortiz-Escobar et al. (2026); Lee et al. (2025); Arystanbekov et al. (2026)
Dependence	New relations of reliance on vendors, connectivity, data flows, subscriptions and update cycles that the person cannot readily see, control or exit	Visibility (whether the person knows what they rely on); exitability (whether an alternative exists); reversibility (whether withdrawal of the system leaves the person worse off than before); concentration (whether reliance is on a single actor)	Cave (2024); Oxley et al. (2025); Ashouri et al. (2026); Wade et al. (2026); Hendy et al. (2026)
Dignity	Whether the arrangement treats the person as an agent whose purposes govern the assistance, and whether the person experiences it that way	Control (who decides what the system does); register (whether the system's language and behaviour are respectful or paternalistic); substitution (whether the system replaces human relations the person values); exposure (whether assistance requires surrendering privacy)	Álvarez-Aguado et al. (2026); Coffey et al. (2026); Bariffi (2026); Soldatic et al. (2026)

The second result is that the three effects are produced by distinct mechanisms and therefore can move independently. Capability extension is driven mainly by model accuracy on the person's own inputs, by latency and by the fit between what the system does and what the person wants to do. Dependence is driven by where computation happens, who owns the model, how it is paid for and

whether it can be replaced. Dignity is driven by how control and inference are distributed, by the priors the model carries about disabled people and by whether assistance displaces relations the person values. Because the drivers differ, a design or provisioning decision can improve one effect while worsening another. On-device inference, for example, reduces dependence and exposure but, with current hardware, constrains model size and so limits extension; cloud inference does the reverse. Personalised speech models increase extension for a single speaker while concentrating dependence on one vendor and making exit costly.

The third result is the set of propositions in Table 2. Each specifies a condition under which the effects align or diverge, with the mechanism and the sources from which it is derived.

Table. 2. Propositions on design and provisioning conditions.

Proposition	Statement	Mechanism	Derived from
P1	Capability extension by AI assistance is largest for the users on whom the underlying model was least trained, and realised least for them	Performance on atypical inputs (impaired speech, blind-taken photographs, atypical movement) degrades sharply; the users with the greatest need face the largest error	Zhao et al. (2025); Tobin et al. (2024); Urqueta Alfaro et al. (2026); Kaboré et al. (2026)
P2	Capability extension is secure, in the capability sense, only when the function survives loss of connectivity, loss of subscription and vendor withdrawal	A functioning contingent on external conditions is an insecure functioning; on-device, open or substitutable systems convert it into a capability	Ortiz-Escobar et al. (2026); Mekonnen et al. (2025); Arystanbekov et al. (2026); Wade et al. (2026)
P3	Dependence rises with the degree to which the system infers rather than executes the person's intention, unless the inference is visible and overridable	Shared autonomy multiplies performance but relocates decision authority to the model; visibility and override preserve the person's agency	Lee et al. (2025); Oxley et al. (2025); Alkhateeb (2026)
P4	Dependence is experienced as diminishing when it is concentrated in a single actor whose continuity the person cannot influence	Implant and personalised-model users are bound to a company's survival and support; aftercare failures are already documented	Ashouri et al. (2026); Oxley et al. (2025); Cave (2024)
P5	Dignity is preserved when the person controls what the system does, and is eroded when the system substitutes for human support the person values	Users of conversational AI with intellectual disability accepted it as a thinking aid and rejected it as a replacement; providers fear loss of relational care	Álvarez-Aguado et al. (2026); Soldatic et al. (2026)
P6	Models carrying paternalistic priors about disabled people will reproduce those priors in assistive roles unless audited for them	Story-generation audits show systematic infantilisation and dependence framing across five large models	Coffey et al. (2026); Dalal et al. (2026)
P7	In safety-critical contexts, probabilistic error converts capability extension into risk, and the conversion is steeper where no human fallback exists	Navigation and communication errors have consequences that accuracy averages conceal; low-resource settings have fewest fallbacks	Singh et al. (2026); Giraldo & Sacchi (2026); Kamenopoulou et al. (2026)
P8	In African provisioning conditions, AI assistance delivered through consumer platforms reaches more people than conventional assistive technology but with higher dependence and lower security, unless the state treats data, language resources and reusable models as public goods	Cost and import dependence block conventional products; smartphones are present but connectivity, electricity, local-language models and repair are not	Pakouyowou (2026); Wade et al. (2026); Acheampong et al. (2026); Zandam et al. (2025); Gaiotto et al. (2026)

A fourth result concerns the risks that AI adds to the older problems of assistive technology provision (cost, abandonment, absent service systems, mismatch with local conditions), which the African

literature reviewed above documents well. AI adds four. The first is probabilistic error: a cane does not hallucinate, but a vision-language model will confidently describe a step that is not there. The second is surveillance: assistance that works by continuously capturing images, audio or neural signals creates a record of the person's life that is held elsewhere, and a subgroup of disabled people in Korea already perceive themselves as more exposed to AI risk than society at large (Kim & Park, 2026). The third is obsolescence of a new kind, in which the device continues to exist but the model behind it is withdrawn, repriced or changed; the "economic dependency" that women in Egypt named among the barriers to AI-supported devices (Hendy et al., 2026) points in the same direction. The fourth is affordability restated: the purchase price of a phone is replaced by a recurring subscription, a data cost and an electricity cost, each a point of failure in low-resource settings.

4.2 DISCUSSION

The framework can be read against the three questions posed in the introduction.

On the first question, the conditions under which AI extends capability, the empirical review supports P1 and P2 in a way that should temper the enthusiasm of the engineering literature. The accuracy gradient documented by Zhao et al. (2025), five per cent error for hearing speakers and over eighty per cent for prelingual deaf speakers, is not a defect to be patched. It follows from the structure of machine learning: a model trained on the distribution of typical bodies and behaviours performs worst on those furthest from it, and those are exactly the people for whom assistance would matter most. Kaboré et al. (2026) show the same structure at the level of research itself, with only six of 93 studies actively involving disabled people in co-design or user validation of the systems built for them. A system that works for a person with mild dysarthria and fails for a person with severe dysarthria has extended capability for the first and, if it is presented as a solution, may have reduced the second person's claim on human support. P2 adds that even where the system works, the extension is only as secure as its conditions. The on-device systems described by Arystanbekov et al. (2026), who run a quantised model with a small loss of benchmark accuracy, and Mekonnen et al. (2025), who run open-source models on a Raspberry Pi, operate without a connection within whatever limits current hardware places on model size, and in doing so convert an insecure functioning into something closer to a capability; that trade-off deserves to be named rather than treated as a technical compromise.

On the second question, when extension becomes dependence, the brain-computer interface literature is the clearest guide because the dependence is physical. Oxley et al. (2025) list continuous and stable use as requirements for autonomy, and Ashouri et al. (2026) point to recent cases of inadequate patient aftercare alongside consumer exploitation and misleading claims. Lee et al. (2025) show the deeper issue. A copilot that infers the user's goal from neural signals, task structure and computer vision and acts on the inference is, in the terms of P3, relocating decision authority. The study reports the gains and counts task errors in the robotic task, and its authors acknowledge that online error correction by the human is harder when the copilot executes a whole action sequence; what it does not report is a separate rate at which the copilot misread what the participant wanted, which is the measure the framework would demand. The point generalises to every predictive interface, from word completion to intention-aware wheelchairs: the more the system does for the person, the more important it becomes that the person can see and refuse what it is doing. Alkhateeb's (2026) finding that metacognitive regulation predicts adaptive reliance suggests the capacity to monitor the system is itself a competence that assistance should build rather than erode.

Cave's (2024) study is the most instructive single case for P4 because it describes a benign vendor. Project Relate is free and was built for exactly this population, and its users liked the concept. Yet their communicative capability became tied to a product whose continuation they could not influence, and a person with amyotrophic lateral sclerosis who trains a personalised model is making an investment that only one company can honour. Concentration of dependence is not a question of intent.

On the third question, dignity, the evidence is thinner and the framework correspondingly more tentative. Álvarez-Aguado et al. (2026) support P5 cleanly: their participants treated the AI as legitimate only when it was accessible, personalised and under their control, valued it for helping them organise their thinking while they retained ownership of the decision, and rejected it as a substitute for human support. Coffey et al. (2026) support P6 by showing that the priors are there to be audited, and Dalal et al. (2026) suggest why they may persist, since most models rely on biomedical and functional data and are trained on datasets that neglect the social and structural determinants of disability. What the

literature does not yet show is how often an assistive system expresses those priors in use, and that is where P6 should be tested first. Soldatic et al. (2026) point to a dimension the framework captures only partly: assistance can be dignified for the person assisted and still unsettle the relations around them, since the care providers in their study feared both for their jobs and for the loss of relational care.

P7 and P8 are where the framework is most exposed and most needed. P7 concerns reliance on a probabilistic system where an error has physical consequences. The engineering literature reports accuracy as an average, and Giraldo and Sacchi (2026) found that only three of sixteen evaluation approaches explicitly evaluated AI-specific properties and that none combined such evaluation with comprehensive coverage of the person, their activities, their participation and their environment. Singh et al. (2026) show the gap from the inside: general-purpose vision-language models convey spatial information poorly, which is exactly what a blind person needs to reach for a hot cup or step off a kerb, and the authors built a dedicated system rather than trust the general model. Whether an error becomes a harm then depends on what else is available; Kamenopoulou et al. (2026) describe deafblind adults in Kenya who relied significantly on human support alongside their assistive products, and where that support is absent and the AI is the sole channel the conversion from error to risk is steepest. No reviewed study measures harm from assistive AI error in use, so P7 rests on the structure of the evidence rather than on incident data. The African studies reviewed here document provision of conventional assistive technology; the one exception, Hendy et al. (2026), reports perceptions of AI-based devices rather than measuring their use, so the propositions about Africa remain inferences from provisioning conditions. Two features of the evidence make the inference defensible. The barriers reported in Ghana, Nigeria, Ethiopia, Kenya, Tanzania and Egypt prominently include cost, maintenance, weak service systems and mismatch with local conditions, alongside inaccessible infrastructure, negative attitudes and socioeconomic disparity (Acheampong et al., 2026; Hendy et al., 2026; Kibret et al., 2025; Kamenopoulou et al., 2026; Ogondiek, 2026; Zandam et al., 2025), and consumer-platform AI changes the form of each barrier without removing it: cost becomes subscription and data, maintenance becomes update dependence, and service becomes a help page. And Wade et al. (2026), in the one recent article that addresses digital assistive technology in low- and middle-income countries directly, argue that affordability, connectivity, electrical reliability and "uneven language infrastructure" determine which digital pathways can be sustained, and propose that data, language resources and foundational tooling be treated as public-interest investments. That is P8 stated as policy. The counter-argument is real: for a blind person in Lagos with a smartphone, an imperfect cloud describer available today may be worth more than a secure device that does not exist, and a framework that discounts insecure extension too heavily would be ableist in its own way. The framework does not rank the effects. It insists that all three be reported.

Three limitations should be stated. The evidence base is skewed towards sensory and motor impairment and towards high-income countries, which is the field's skew rather than this paper's but constrains what can be claimed. The propositions are derived from convergence across small and heterogeneous studies, several of which are reviews rather than primary research, and should be read as hypotheses. And the framework treats the person as the unit of analysis, which under-represents the family and community relations through which assistive technology is actually used in most African settings.

Future research follows from the propositions. P1 could be tested by reporting error rates stratified by impairment severity as a standard element of every assistive AI evaluation. P3 could be tested by measuring override rates and inference errors in shared-autonomy interfaces. P5 and P6 call for audits of assistive agents in use, with disabled users rating register and control. P7 calls for incident reporting: how often an assistive error in navigation or communication led to harm, and whether a human fallback was available at the time. P8 requires the first studies of AI-based assistance as it is actually used in African cities, which the provisioning literature is well placed to conduct if it extends its instruments to functions delivered through phones.

5. CONCLUSION AND RECOMMENDATION

The paper set out to say what "works" should mean for an AI-based assistive technology and has argued that it must mean three things at once. The system should enlarge the real opportunities of the person using it, and do so securely. It should not bind the person to actors and infrastructures they cannot see, influence or leave. And it should leave the person in command of the assistance and treated as the agent whose purposes it serves. The empirical literature shows that these effects are produced by different

mechanisms, routinely diverge, and diverge most sharply for the people with the greatest need and in the settings with the weakest provision. Evaluation that reports only the first effect is not neutral; it systematically favours systems that are powerful and fragile over systems that are modest and secure.

Several recommendations follow. Ministries of health and social development that are developing national assistive technology strategies, priority product lists and coordination mechanisms should extend them to AI-delivered functions, specifying for each the conditions under which it is secure: offline operation, local-language support, substitutability and a named service pathway. Data protection authorities and sector regulators should require that assistive AI systems declare what data leaves the device, who holds it and for how long, and should treat the silent withdrawal or repricing of an assistive model as a consumer harm. Public procurers should weight security of functioning and exitability alongside accuracy. Universities and research funders in Africa should invest in the public goods Wade et al. (2026) identify, namely representative datasets and language resources, and reusable models and reference toolchains that local firms can build on, because without them the accuracy gradient of P1 will be steepest on the continent. Designers should report error rates by impairment severity, make inference visible and overridable, and audit conversational agents for paternalistic register before deployment. Organisations of persons with disabilities should be involved before the design is fixed rather than at the testing stage, and resourced to run their own audits of systems marketed to their members. None of this requires new technology. Each requires that the three effects be named and counted.

6. CONTRIBUTION TO EXISTING KNOWLEDGE

The paper's first contribution is conceptual. It separates three effects that the literature on AI and disability has treated as a single question of benefit, defines each in terms that can be operationalised, and shows through the capability approach why they can move in opposite directions. The distinction between extending a functioning and securing a capability, and the treatment of dependence as a matter of visibility, exitability, reversibility and concentration rather than of presence or absence, give evaluators a vocabulary that accuracy metrics and satisfaction scales lack.

The second contribution is propositional. Eight propositions, each grounded in at least two independent recent sources, specify conditions under which capability extension coincides with autonomy and dignity and conditions under which it does not. They are phrased to be tested and several indicate the measure that would test them.

The third contribution is to bring two literatures into contact. The engineering and human–AI interaction literature on assistive AI and the global assistive technology provisioning literature describe the same people and seldom cite each other. By reading the African provisioning studies against the AI studies, the paper identifies the distinctive risks AI adds to old provisioning problems and argues that the continent is the framework's most demanding test rather than a peripheral case. The propositions about African conditions are the least evidenced in the paper and are offered as an agenda for the provisioning literature to take up.

7. REFERENCES

Acheampong, E., Kasore, V., Kponyo, B. A., Owusu-Ansah, R. Y., Owusu-Ansah, J. K., Ansah, O., Owusu Ansah, E., Owusu Agyemang, J., Mensah Gyening, R.-M. O., Gyening, M., Ahene, E., & Kponyo, J. J. (2026). Barriers to assistive technology uptake among persons with disabilities in selected urban districts in Ghana. *Scientific Reports*, *16*(1), 1838. https://doi.org/10.1038/s41598-025-31466-4

Alkhateeb, N. (2026). Generative AI as a cognitive assistive technology for adults with dyslexia: The role of anxiety, working memory, and metacognitive regulation. *Disability and Rehabilitation: Assistive Technology*. Advance online publication. https://doi.org/10.1080/17483107.2026.2715246

Alsaleh, A. (2025). The influence of artificial intelligence on individuals with disabilities. *Acta Psychologica*, *262*, 106010. https://doi.org/10.1016/j.actpsy.2025.106010

Álvarez-Aguado, I., Vega, V., Rosselló-Peñaloza, M., Aguado, L. Á., Vilo, K. C., Figueroa, J. F., La Rivera, F. M., González-Carrasco, F., & Parra, F. E. (2026). Conversational AI as personalised cognitive support: Adults with intellectual disability's experiences of everyday decision-making. *Journal of Intellectual Disability Research*. Advance

online publication. https://doi.org/10.1111/jir.70171

Anderson, S., Seghers, F., Saunders, L., Reza, C., Chikhata, F., Bekele, E., Okiring, I., Mugambi, K., Alesinloye, G. I., & Mishra, S. (2026a). Building effective leadership and coordination for assistive technology: Best practices and recommendations from six countries across Sub-Saharan Africa. *Disability and Rehabilitation: Assistive Technology*. Advance online publication. https://doi.org/10.1080/17483107.2026.2717494

Anderson, S., Seghers, F., Saunders, L., Reza, C., Chikhata, F., Bekele, E., & Mishra, S. (2026b). Developing National Priority Assistive Products Lists: Experiences and lessons from six Sub-Saharan African countries. *Disability and Rehabilitation: Assistive Technology*, *21*(2), 668–684. https://doi.org/10.1080/17483107.2025.2561246

Arystanbekov, B., Varol, H. A., & Yazici, A. (2026). Edge-based vision-language assistive system for the visually impaired: A quantized VLM approach. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. Advance online publication. https://doi.org/10.1109/TNSRE.2026.3696447

Ashouri, D., Weh, L., Borrmann, V., Coenen, C., Kraft, E. M., Loufs, V., Mehnert, W., Muhr, P., Müller, O., Şahinol, M., Schmidt, M., Seyfried, G., Soedibjo, K., Stewart, I. E., Wolbring, G., Youssef, S., & Stieglitz, T. (2026). Neurotechnology through the lens of users and as a novel field for society. *Frontiers in Neurology*, *17*, 1842746. https://doi.org/10.3389/fneur.2026.1842746

Bariffi, F. (2026). Mind, machine, and the law: Reimagining neurotechnology governance through disability rights. *Journal of Law and the Biosciences*, *13*(1), lsag011. https://doi.org/10.1093/jlb/lsag011

Cave, R. (2024). How people living with amyotrophic lateral sclerosis use personalized automatic speech recognition technology to support communication. *Journal of Speech, Language, and Hearing Research*, *67*(11), 4186–4202. https://doi.org/10.1044/2024_JSLHR-24-00097

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*, *19*(4), 102086. https://doi.org/10.1016/j.dhjo.2026.102086

Curati-Alasonatti, E., Layton, N., Kashyap, A., & Mishra, S. (2026). Lessons learnt from assistive technology ecosystem strengthening programmes in 12 low- and middle-income countries. *Disability and Rehabilitation: Assistive Technology*, *21*(6), 3303–3316. https://doi.org/10.1080/17483107.2026.2712129

Dalal, J., Teodoro, D., & Rivas Velarde, M. (2026). The promise and peril of AI for disability health equity: A justice-centered narrative review. *PLOS Digital Health*, *5*(9), e0001663. https://doi.org/10.1371/journal.pdig.0001663

Du, W., Peng, X.-Y., Li, W., Attiq, U.-R., Li, X.-M., Feng, R., Zhong, K., Lu, X.-R., & Chen, H.-L. (2026). Mapping applications of artificial intelligence in social support for persons with disabilities: A systematic scoping review. *Rehabilitation Psychology*. Advance online publication. https://doi.org/10.1037/rep0000692

Gaiotto, E. M. G., Gabriel, M., Fernandes, M. F. P., de Andrade, O. G., Hedlund Jardim, N., Ramos, V. D., Souza, N. M., & Aguiar, B. G. A. (2026). Access to assistive technology in Africa: Findings from evidence-informed deliberative dialogues on trade and fiscal policies. *Frontiers in Rehabilitation Sciences*, *7*, 1866962. https://doi.org/10.3389/fresc.2026.1866962

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

Hendy, A. A., Rehim, M., Kaouache, R., & Yaiche, O. (2026). AI and digital assistive technologies and reproductive rights among women with disabilities in Egypt: A qualitative SWOT study. *Frontiers in Sociology*, *11*, 1925235. https://doi.org/10.3389/fsoc.2026.1925235

Joskow, R., Patel, D., Landre, A., Mattick, K., Holloway, C., Danemayer, J., & Austin, V.

(2025). Understanding the impact of assistive technology on users' lives in England: A capability approach. *Bioengineering*, *12*(7), 750. https://doi.org/10.3390/bioengineering12070750

Kaboré, S. S., Bayani, A., Ewane, L. P. E., Dos Anjos, D. S. O., Nikiéma, J. N., & Mac-Seing, M. (2026). Contribution of artificial intelligence research to the development of inclusive applications for people with disabilities: A scoping review (2015–2025). *PLOS Digital Health*, *5*(10), e0001749. https://doi.org/10.1371/journal.pdig.0001749

Kamenopoulou, L., Prain, M., Bell, D., Maragia, M., Ranzato, E., & Layton, N. (2026). Exploring the assistive technology needs and rights of four Kenyan adults with deafblindness by using the My Assistive Technology Outcomes Framework. *Disability and Rehabilitation: Assistive Technology*, *21*(7), 3639–3657. https://doi.org/10.1080/17483107.2026.2664019

Kibret, A. K., Yilma, T. M., Davies, T. C., Alemayehu Genetu, Z., Mekonone, S. T., Alene, B. M., Addis, B. M., Nigussie, K., Ayele, R. G., & Mekonnen Araya, B. (2025). Assistive technology and daily living challenges among students with disabilities at University of Gondar, Ethiopia: A qualitative study. *Disability and Rehabilitation: Assistive Technology*, *20*(1), 201–210. https://doi.org/10.1080/17483107.2024.2362248

Kim, J., & Park, E. (2026). When the risk feels personal: Self-relative pessimism in AI-related risk perception among people with disabilities. *Disability and Rehabilitation: Assistive Technology*. Advance online publication. https://doi.org/10.1080/17483107.2026.2719080

Lee, J. Y., Lee, S., Mishra, A., Yan, X., McMahan, B., Gaisford, B., Kobashigawa, C., Qu, M., Xie, C., & Kao, J. C. (2025). Brain–computer interface control with artificial intelligence copilots. *Nature Machine Intelligence*, *7*(9), 1510–1523. https://doi.org/10.1038/s42256-025-01090-y

Mahmoudi, E. (2026). Techno-solutionism in the age of artificial intelligence: A biopsychosocialtech perspective for assistive technology. *Disability and Rehabilitation: Assistive Technology*, *21*(6), 3330–3331. https://doi.org/10.1080/17483107.2026.2648859

Marul, E. (2026). Assistive communication technologies for individuals with multiple disabilities: An evidence mapping review informing future AI-enabled AAC. *Disability and Rehabilitation: Assistive Technology*, *21*(7), 3537–3558. https://doi.org/10.1080/17483107.2026.2682413

Mekonnen, F. Y., Al Bataineh, M. F., Abu Abdoun, D., Serag, A., Tamiru, K. T., Abula, W., & Darota, S. (2025). Development of a fully autonomous offline assistive system for visually impaired individuals: A privacy-first approach. *Sensors*, *25*(19), 6006. https://doi.org/10.3390/s25196006

Ogondiek, J. W. (2026). Policy-practice gap: A document analysis of assistive technology access for pupils with disabilities in Tanzania. *Disability and Rehabilitation: Assistive Technology*, *21*(7), 3559–3572. https://doi.org/10.1080/17483107.2026.2683132

Ortiz-Escobar, L. M., Chavarria, M. A., Hurst-Majno, S., Campo Salazar, O. I., Escobar-Hurtado, C., Stein, M. A., & Rivas Velarde, M. (2026). A social justice approach to assistive technology and well-being of people with visual disabilities in low- and middle-income countries: Qualitative narrative study. *JMIR Rehabilitation and Assistive Technologies*, *13*, e72306. https://doi.org/10.2196/72306

Oxley, T. J., Deo, D. R., Cernera, S., Sawyer, A., Putrino, D., Ramsey, N. F., & Fry, A. (2025). The 'Brussels 4': Essential requirements for implantable brain–computer interface user autonomy. *Journal of Neural Engineering*, *22*(1), 013002. https://doi.org/10.1088/1741-2552/ada0e6

Pakouyowou, M. T. (2026). From aid dependency to market transformation: Harnessing international trade and innovative business models to revitalize the assistive technology ecosystem in Africa. *Frontiers in*

Rehabilitation Sciences, *7*, 1905487. https://doi.org/10.3389/fresc.2026.1905487

Ryan, S., & O'Brien, R. (2024). Using a capability approach to explore how people with intellectual disabilities can lead flourishing lives. *Journal of Applied Research in Intellectual Disabilities*, *37*(6), e13299. https://doi.org/10.1111/jar.13299

Singh, A., Bhanushali, M. A., Luo, J., Luo, G., & Pundlik, S. (2026). Assisting the blind to reach daily objects using smart glasses. *Displays*, *93*, 103390. https://doi.org/10.1016/j.displa.2026.103390

Soldatic, K., Balram, R., Lee, M., Santilli, T., & Magee, L. (2026). Relational meanings of AI in disability care: An intersectional, arts-based inquiry. *Qualitative Health Research*, *36*(2-3), 304–316. https://doi.org/10.1177/10497323251401541

Tobin, J., Nelson, P., MacDonald, B., Heywood, R., Cave, R., Seaver, K., Desjardins, A., Jiang, P.-P., & Green, J. R. (2024). Automatic speech recognition of conversational speech in individuals with disordered speech. *Journal of Speech, Language, and Hearing Research*, *67*(11), 4176–4185. https://doi.org/10.1044/2024_JSLHR-24-00045

Umbelino, L. L., Ferreira, D. M. T. P., & Alonso, C. M. do C. (2026). Ergonomics in assistive technology design: Scoping review. *Disability and Rehabilitation: Assistive Technology*, *21*(7), 3469–3489. https://doi.org/10.1080/17483107.2026.2671838

Urqueta Alfaro, A., Roundtree, K. A., Pfaff, M. S., Bos, N., Edell, J. D., & ten-Brink, R. (2026). Automatic speech recognition underperforms with diverse accents as measured by word error rate and a semantic similarity measure. *Disability and Rehabilitation: Assistive Technology*. Advance online publication. https://doi.org/10.1080/17483107.2026.2716598

Wade, W., Hu, S., & Fleming, J. (2026). Market shaping via digital public infrastructure for assistive technology in low- and middle-income countries. *Frontiers in Rehabilitation Sciences*, *7*, 1881812. https://doi.org/10.3389/fresc.2026.1881812

Wilkens, J., Banerji, R., & Mishra, S. (2026). Financing assistive technology: Policy options on the path towards universal health coverage. *Frontiers in Public Health*, *14*, 1901297. https://doi.org/10.3389/fpubh.2026.1901297

Zandam, H., Sulaiman, S. K., Mohammad, A. H., & Sabo, A. (2025). Utilization and impact of assistive device use on mental health among children with disabilities in Nigeria: A cross-sectional study. *Disability and Rehabilitation: Assistive Technology*, *20*(7), 2094–2102. https://doi.org/10.1080/17483107.2025.2524062

Zhao, R., Choi, A. S. G., Koenecke, A., & Rameau, A. (2025). Quantification of automatic speech recognition system performance on d/Deaf and hard of hearing speech. *The Laryngoscope*, *135*(1), 191–197. https://doi.org/10.1002/lary.31713