IIAR JOURNALS

Innovating Knowledge. Advancing Research.

IIAR Journal of Disability, Inclusion, and Assistive Technologies

Vol 1, No 1, 2026

Inclusive Education in the Age of Generative AI: A Theoretical Framework of Personalisation, Accommodation and Substitution for Students with Disabilities

IGWE Emmanuel Sunday[1], NWOGBOJI Emmanuel Usulor[2]

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

ABSTRACT

Generative artificial intelligence (AI) can now produce, on demand, what inclusive education has promised for three decades and seldom delivered: material rewritten for a particular reader, explanations in the form a particular student needs, and reading and writing support available at any hour without a specialist in the room. The question this paper asks is under what conditions that capacity produces inclusion rather than a quieter form of segregation. The problem is that the literatures on generative AI in education and on inclusive education have developed apart: the first measures learning gains and dependence, the second asks about presence, participation and belonging, and neither can say when an AI-supported student has been included and when the task has simply been taken away from them. The paper is a theoretical synthesis. Drawing on inclusive pedagogy, universal design for learning, self-determination theory, the social model of disability and the accommodation logic of educational measurement, and on peer-reviewed research published between 2023 and 2026 on generative AI in education, on students with disabilities and on inclusive education in African school systems, it distinguishes three things an AI system can do for a disabled learner: personalise content, provide accommodation, and substitute for effort. The first two advance inclusion; the third undermines it, because an accommodation is meant to remove a barrier to the task rather than the task itself, and the line between the two is drawn by the learning objective, not by the tool. The framework is set out in nine propositions covering the classification of AI uses, the shifting roles of teachers, special educators and peers, the system-level risk that AI accommodation becomes a cheap substitute for human support and physical accessibility, the conditions under which learner agency survives, and the infrastructural and linguistic conditions that govern whether any of this reaches African classrooms. Recommendations are made for ministries of education, examination bodies, teacher-education institutions, school leaders, disabled people's organisations and researchers, with particular attention to school systems in which specialist support is scarcest.

Keywords: inclusive education; students with disabilities; generative AI; personalised learning; accommodation

Inclusive Education in the Age of Generative AI: A Theoretical Framework of Personalisation, Accommodation and Substitution for Students with Disabilities

1. INTRODUCTION

A student with dyslexia in a secondary school in Ibadan has, in principle, access to the same writing assistant as a student in Oslo. For the first time, the gap between what inclusive education recommends and what a teacher with forty-five pupils can actually do is not primarily a gap of labour: a text can be simplified, re-levelled, read aloud, translated or turned into questions in the time it takes to type a request. The promise is obvious enough to have been made many times already, usually by people selling something. The problem for inclusive education is more awkward. The same tool that removes a barrier to a task can remove the task, and from the outside the two look identical: a student who uses a language model to fix the spelling in an essay they composed and a student who has had the essay composed for them both hand in a well-spelled essay.

This paper is about that distinction and what follows from it, and it starts from a position that is not neutral. Inclusive education, as argued for since the Salamanca Statement and codified in Article 24 of the Convention on the Rights of Persons with Disabilities, is a claim about participation in a common curriculum alongside peers, not a claim about outputs. A technology that lets a student produce the right output while participating less in the common work of the classroom is not, on this reading, inclusive, however well the student scores. Equally, a school that forbids the technology in the name of academic integrity, removing the only means by which some students can take part at all, is not protecting standards; it is re-segregating by other means. Both failures are already visible, and both will be more common where specialist support is scarcest, which is to say in most African classrooms.

The research field is young and uneven. Within three years of the public release of large language models the education literature had produced reviews of the opportunities and challenges of AI in education generally (Chiu et al., 2023) and of large language models in particular (Kasneci et al., 2023), a meta-analysis of effects on engagement (Xia et al., 2025) and a growing experimental literature on what unrestricted access does to learning (Bastani et al., 2025; Fan et al., 2025; Stadler et al., 2024). The literature on students with disabilities has lagged. The systematic review of AI in special education by Hopcan et al. (2023) predates generative tools almost entirely; a scoping review of generative AI for neurodivergent school students found only nine sources with original data (Ronksley-Pavia et al., 2025); and a bibliometric review of generative AI in disability-inclusive learning found research concentrated on autism and dyslexia, with visual, hearing and motor impairments under-represented (Amal Krishnan et al., 2026). What exists is largely descriptive, short-term and conducted in well-resourced systems, and the two literatures do not share a vocabulary: the generative-AI literature speaks of over-reliance, offloading and agency, the inclusive-education literature of presence, participation and belonging, and neither can say on its own when an AI-supported student has been included.

The purpose of this paper is to supply that vocabulary. It distinguishes personalisation, accommodation and substitution as three things a generative-AI system can do for a disabled learner, argues that the first two advance inclusion while the third undermines it, and shows that the boundary between accommodation and substitution is a property not of the tool but of the learning objective against which it is used. From that distinction it derives propositions about the changing work of teachers, special educators and peers, about the system-level risk that AI accommodation becomes a cheaper substitute for human support and physical accessibility, about the conditions under which learner agency survives, and about the infrastructural and linguistic conditions that decide whether any of this reaches the students who most need it. The paper is theoretical and non-empirical. Three questions organise it: what does generative AI do for a disabled learner that is different from what it does for any learner; what distinguishes an AI use that includes from one that merely produces; and what conditions must hold, in the classroom and in the system, for the inclusive uses to dominate.

2. LITERATURE REVIEW

2.1 CONCEPTUAL REVIEW

Three concepts carry the argument. The first is inclusion itself, a term stretched to cover everything

from the physical placement of disabled children in mainstream schools to whole-school reform of culture and curriculum, and the stretching matters because what counts as success depends on which meaning is in play: placement-based definitions are satisfied when the child is in the room, participation-based definitions ask whether the child is doing the same work with the same people, and outcome-based definitions ask whether the child learns as much. Studies of inclusive practice in African systems show how far apart these can sit: Maphie (2025) found Tanzanian students with disabilities placed in general classrooms or in classrooms paired with a resource room, with support described as emotional, spiritual, social, material and academic, and Chirowamhangu (2024) found South African provision still marked by shortages of specialist schools, transport, expert support and basic infrastructure two decades after the policy meant to end them. This paper uses a participation-based definition, and it does so deliberately, because the output-producing power of generative AI makes outcome-based definitions dangerously easy to satisfy.

The second concept is accommodation, which in educational measurement has a precise meaning the wider literature tends to blur. An accommodation changes how a student accesses a task or demonstrates what they know without changing what the task measures. Extra time, a reader, a scribe and large print are accommodations because they address construct-irrelevant barriers: a reading disability is irrelevant to the construct in a mathematics examination, so a reader does not change what is measured. A modification changes the construct; a simplified mathematics test measures something different. The distinction is the only principled way to say when assistance stops being access, and generative AI makes it harder to apply for two reasons. A language model does not provide a single, bounded accommodation; the same interface will read a text aloud, summarise it, explain it and write the answer, and the student chooses where to stop. And for many disabilities the barrier and the target skill are entangled: a reader is a clean accommodation for a student with dyslexia in a history examination, but in a writing class the student's difficulty with text is the very thing being taught.

The third concept is agency, and here the generative-AI literature has done most of the work. Darvishi et al. (2024) showed in a large randomised study that students given AI feedback assistance relied on it, that performance fell back when it was withdrawn, and that prompts to self-regulate only partly protected them; Fan et al. (2025) found that students using ChatGPT improved essay scores without improving knowledge gain or transfer, and named the pattern metacognitive laziness; Zhai et al. (2024), reviewing fourteen studies of over-reliance on AI dialogue systems, found that such reliance affected students' decision-making, critical thinking and analytical reasoning. Agency in this paper means something narrower than these studies sometimes imply: the learner's control over what the tool does, their understanding of what it has done, and their retained capacity to do the task when the tool is absent. The last, which Chen (2026) separated as supported performance versus independent no-AI outcomes, is the one inclusive education cannot afford to lose sight of.

2.2 THEORETICAL REVIEW

Five bodies of theory inform the framework, and they pull in different directions, which is part of the point. The accommodation logic of educational measurement, already set out in Section 2.1, supplies the test for where assistance stops being access; the other four are taken in turn here. Inclusive pedagogy rests on extending what is ordinarily available to all learners rather than providing something additional or different for some, as a deliberate response to the segregating effect of individualised provision: a child who is always given different work is marked out by it. Universal design for learning (UDL) makes a related argument at the level of curriculum design, holding that materials should offer multiple means of representation, engagement and expression from the outset so that fewer learners need retrofitted accommodation. Bray et al. (2024), reviewing fifteen studies of technology in UDL implementations at second level, found the research concentrated on the representation principle, mainly by giving learners options for accessing content, with a clear gap on technology for engagement and for action and expression. Both traditions are suspicious of personalisation when it means giving each child something different, and both would regard an AI that quietly re-levels every text for every child as a return to the watered-down curriculum inclusive education was meant to end. Hyatt and Owenz (2026) paired an AI-based assignment with UDL options for demonstrating learning in graduate classes with a high proportion of students with disabilities, and McCree et al. (2026) offer special educators a framework for planning AI-supported instruction under UDL and culturally responsive teaching; the question this paper raises is what stops cheap

variation of representation sliding into variation of the objective.

Self-determination theory supplies the account of agency: learning that is autonomously motivated, in which the learner experiences competence, autonomy and relatedness, is more durable and transferable than learning that is controlled. Fan et al. (2026), synthesising forty-six primary studies of generative AI use through that lens, concluded that outcomes depended less on how often students used AI than on whether they retained cognitive agency and regulated when and why they delegated, and that controlling or weakly structured environments promoted dependent offloading. The theory's emphasis on relatedness also warns against support that isolates: a student whose every need is met by a private tool has gained autonomy from peers and lost the relationships through which belonging is built.

The last is the social model of disability, which locates disability in the mismatch between a person and an environment rather than in the person, and on which generative AI is simply another feature of the environment that can lower or raise barriers depending on how it is deployed. The model is often read as uniformly hopeful about technology, but it carries a warning that matters here. If the environment is changed by giving the disabled student a tool while leaving the classroom, the curriculum and the assessment untouched, the barrier has been privatised rather than removed; the student now carries the adjustment as a possession that may be withdrawn, may fail, and may be seen by others as an advantage rather than an entitlement. Nor is the tool neutral about disability: Urbina et al. (2025) found that ChatGPT and Gemini both described people with disabilities with fewer favourable and more limiting qualities than non-disabled people, and Hadar Souval et al. (2025) caught the doubleness in their image of generative AI as at once a social mirror of what society thinks about cognitive disability and a cognitive copilot for the people it describes so poorly.

2.3 EMPIRICAL REVIEW

The empirical literature falls into three uneven piles. The largest concerns generative AI and learning in general, and its central finding is that unrestricted access tends to raise performance on the supported task and lower, or leave unchanged, performance when support is removed. Bastani et al. (2025), in a field experiment with nearly a thousand Turkish secondary-school students, found that a standard ChatGPT interface improved practice performance substantially but lowered subsequent unassisted examination scores by about seventeen per cent, whereas an interface with tutoring guardrails raised practice performance without the later cost. Stadler et al. (2024) found that students using a language model for scientific inquiry experienced less mental effort but reasoned more shallowly. Chen (2026) found the same divergence in an eight-week writing study: open AI collaboration produced the best supported essays, while a bounded condition with compulsory reflection produced better independent writing, higher-order thinking and argument depth afterwards. Fan and Chang (2026), randomising over seven hundred undergraduates to different divisions of labour between student and AI in report writing, found that arrangements in which the student judged AI-produced text outperformed those in which the AI judged the student's text, and, more tellingly, that students with a high tendency to offload re-outsourced the evaluation they had been assigned. Their summary, that prescribing a division of labour governs the assignment but not the cognition, is a sentence every accommodation policy should have pinned above it. Choudhary et al. (2026) added a mechanism: perceived sycophancy in AI tutors predicted metacognitive laziness, which predicted dependence, which predicted lower autonomy. Against this, Xia et al. (2025) found in a meta-analysis of seventeen studies that ChatGPT-based learning improved behavioural, cognitive and emotional engagement relative to learning without it; it would be a misreading of the dependence literature to conclude that the tools do harm as such, since the consistent message is that design and framing decide the effect.

The second pile concerns students with disabilities, and it is thinner and more optimistic. Gharaibeh and Basulayyim (2025) reported in a quasi-experiment with sixty Arabic-speaking children with dyslexia that ChatGPT-based reading instruction improved sentence comprehension relative to controls. Alkhateeb (2026), following over five hundred adults with dyslexia across two waves, found that anxiety predicted declines in metacognitive regulation and in adaptive reliance on AI, while gains in regulation predicted more adaptive reliance; the capacity to use AI well is itself unevenly distributed, and trainable. Elshaer et al. (2025) found among 385 visually impaired university students that trust in large language models predicted use, which improved quality of life, with academic success affected only indirectly. Flütsch Keravec et al. (2026) asked, in an intervention study, whether composing by speech-to-text changes writing acquisition for students with dyslexia, which is precisely the question of whether a

tool that bypasses transcription supports or supplants the target skill. For deaf and hard-of-hearing students, Alkahtani (2024) synthesised twenty studies and found that automatic speech recognition, AI-enabled captioning and related tools increased classroom participation by removing barriers, while Jolly et al. (2024) showed from the narratives of fifteen deaf college students that many eventually gained effective access through real-time captioning but could initially struggle to overcome barriers to using the service. Seung and Basham (2026) framed the question for students with learning disabilities directly as one of cognitive offloading, asking what the general findings mean for a population whose disabilities often lie in the very processes being offloaded. On the teacher side, Rakap and Balikci (2026) found in a small randomised study that special educators using ChatGPT wrote better individualised education plan goals; van Garderen et al. (2025) showed how generative AI can produce scaffolded resources for students with disabilities; and Zraydi et al. (2026) found among 244 general and special education teachers in the United Arab Emirates that AI-related technological pedagogical content knowledge was related to intentions to use AI for students with learning disabilities, with only one of its components predicting those intentions. Morgan (2026) argued from UK equality law that regulated use of a language model may be a reasonable adjustment for a neurodivergent student when the barrier lies in written expression rather than reasoning, which requires universities to distinguish the competence a course assesses from the method by which it assesses it.

The third pile concerns African education systems, where the generative-AI literature almost disappears and what remains is a literature on inclusive education under constraint; the AI evidence that does exist comes from universities and from outside the continent, and its application to schools is an extrapolation the paper makes openly. Motitswe (2025) found South African school-based support teams reporting inadequate training, overcrowded classrooms, weak policy guidance and a persistent medical model of disability; Malahlela and Johnson (2024) found teachers who understood inclusion policy in theory but had little confidence applying it; Kuyini (2026), surveying 448 principals and teachers in sixty Ghanaian primary schools, found principals holding only low to moderate expectations that teachers would practise inclusively at all. On AI, Sharma and Ha (2025) found ChatGPT adoption high among students in both the United States and Bangladesh, predicted by social persuasion, perceived usefulness and country rather than individual characteristics, and suggested the tools may matter more for learning in the Global South; Sevhake and Hofisi (2026) examined through document analysis how South African and Kenyan universities are weighing efficiency against equity and student autonomy. Nothing in the peer-reviewed literature to 2026 reports a controlled study of generative AI with disabled learners in an African school. That absence is itself a finding, and the framework treats it as one.

3. MATERIALS AND METHODS

This is a theoretical paper and its method is conceptual synthesis: the construction of a framework from existing theory and evidence, the derivation of propositions from that framework, and the examination of those propositions against the empirical literature. No human participants were involved, no primary data were collected and no ethics approval was required. The only dataset is the reference list, which is openly available with digital object identifiers.

The literature was assembled in October 2026 and restricted by design to peer-reviewed journal articles published between January 2023 and October 2026, the period in which generative AI has been available to students and teachers. Preprints, conference proceedings, book chapters and reports were excluded; the Convention on the Rights of Persons with Disabilities is referred to as a legal instrument, not as evidence. The Crossref bibliographic index and PubMed were queried with combinations of terms for generative AI (generative AI, ChatGPT, large language model), disability (disability, special educational needs, neurodivergent, dyslexia, autism, visual impairment, deaf), inclusion (inclusive education, universal design for learning, accommodation, reasonable adjustment), agency (learner agency, self-regulated learning, cognitive offloading, metacognitive) and region (Africa, Nigeria, Ghana, Kenya, South Africa, Tanzania); abstracts were retrieved through publisher pages, PubMed and Semantic Scholar. Approximately 240 records were screened on title, about 90 were read at abstract level, and 40 are cited.

Inclusion required original research, a systematic or scoping review, or developed theory bearing on at least one of the three functions, in a journal article with a Crossref record inside the date window; articles were excluded if they concerned AI in education without bearing on disability, inclusion, learner agency or the policy and teacher-competence

conditions under which the tools are used, if their venue could not be confirmed as a peer-reviewed journal, or if they duplicated a stronger study. Every cited article was verified against its Crossref publisher record for type, authors, title, journal, year, volume, issue and pagination, and each record was checked for retraction, correction or expression-of-concern notices. One cited article (Bastani et al., 2025) carries a published correction that does not affect the result relied on here; no cited article is retracted or under concern. Where a journal's print issue year differs from the online year, the issue year is used.

The analysis proceeded in three steps. The three functions (personalisation, accommodation, substitution) were first defined by reference to inclusive pedagogy, UDL, self-determination theory, the social model of disability and the accommodation logic of educational measurement, and tested against examples from the empirical literature for mutual exclusivity relative to a given learning objective, since the same action can fall under different functions when the objective changes. Propositions were then derived by asking, for each construct, what conditions in the classroom and in the system would make it more or less likely to dominate, and what it implies for the roles of teacher, special educator and peer. Each proposition was finally examined against the empirical literature to identify where it is supported, where the evidence is thin and where it is contradicted (Section 4.2). The basis of judgement throughout is explicit: a use of AI counts as inclusive if it increases the student's participation in the common curriculum alongside peers without reducing their independent capacity to do what the curriculum sets out to teach.

An AI writing assistant was used during the preparation of this manuscript to support literature retrieval, drafting and formatting, under the direction of the author, who selected the sources, verified every reference against its publisher record, determined the argument and takes full responsibility for the content. No AI system was used to generate or analyse data.

4. RESULTS AND DISCUSSION

4.1 RESULTS

The framework rests on a single distinction, set out in Table 1, between three things a generative-AI system can do for a disabled learner. The distinction is not a taxonomy of tools. The same chat interface can do all three in the same session, and the same action can fall under different headings depending on what the student is supposed to be learning.

Table. 1. The three functions of generative AI for a disabled learner

Function	What the system does	Relation to the learning objective	Inclusive-education reading	Illustrative evidence
Personalisation	Varies the representation, pace, examples or language of content for a particular learner	Objective held constant; the route to it varies	Advances inclusion when it is available to all and does not vary the objective; re-segregates when it quietly lowers the objective for some	Bray et al. (2024); van Garderen et al. (2025); Gharaibeh & Basulayyim (2025)
Accommodation	Removes a barrier that is irrelevant to the objective (reads text aloud, transcribes speech, captions, reorganises, reminds)	Objective untouched; a construct-irrelevant barrier is removed	Advances inclusion; this is what reasonable adjustment has always meant	Alkahtani (2024); Jolly et al. (2024); Morgan (2026); Flütsch Keravec et al. (2026) (open question: accommodation or substitution for the skill being taught)
Substitution	Performs the task, or the part of it, that the objective is meant to develop	Objective bypassed; the construct-relevant work is done by the system	Undermines inclusion by producing presence and output without participation or learning	Bastani et al. (2025); Fan et al. (2025); Chen (2026); Seung & Basham (2026)

Two features of Table 1 do the work. The boundary between accommodation and substitution runs through the learning objective, not the tool: a system that drafts a paragraph is an accommodation for a student with a motor impairment in a history class and substitution in a composition class, and nothing about the software tells the student which situation they are in. And personalisation sits on both sides of the line: it advances inclusion when it varies the route to a common objective and is available to every learner, and re-segregates when it varies the objective for some, more quietly than the old practice of different worksheets because the re-levelled text looks like the original.

From that distinction nine propositions follow (Table 2).

Table. 2. Propositions of the personalisation–accommodation–substitution framework

No.	Proposition	Mechanism	Principal supporting literature
P1	AI personalisation advances inclusion when it varies representation, pace and means of expression while holding the learning objective constant; when it varies the objective it reproduces segregation in a less visible form	Inclusive pedagogy's distinction between extending what is ordinarily available and providing something different; UDL's multiple means	Bray et al. (2024); Hyatt & Owenz (2026); McCree et al. (2026)
P2	The same AI action is an accommodation in one task and a substitution in another; the classification requires an explicit statement of which skill is the target of the task and which is an access skill	Accommodation logic of measurement: construct-relevant versus construct-irrelevant barriers	Morgan (2026); Fan & Chang (2026)
P3	Substitution risk is higher for students with disabilities than for other students, because the barrier and the target skill are more often entangled and because the effort the tool removes is larger	Entanglement of impairment and target skill (dyslexia and writing; executive-function difficulty and planning); larger effort differential	Seung & Basham (2026); Alkhateeb (2026); Flütsch Keravec et al. (2026) (open question)
P4	The teacher's role shifts from delivering content to specifying objectives and adjudicating the accommodation–substitution boundary; the special educator's role shifts from producing adapted material to designing, monitoring and fading AI-supported accommodation	Objective specification becomes the scarce input once material production is cheap	Marino et al. (2023); Rakap & Balikci (2026); van Garderen et al. (2025); Zraydi et al. (2026)
P5	AI accommodation that privatises support reduces the learner's dependence on peers and, with it, the interactions through which belonging is built; inclusion is relational, so presence without participation is possible	Self-determination theory's relatedness need; participation-based definition of inclusion	Fan et al. (2026); Jolly et al. (2024)
P6	Where specialist support, assistive devices and accessible buildings are scarce, AI accommodation will be adopted as the cheaper substitute for them, and the risk of system-level substitution is greatest where support is scarcest	Cost asymmetry between a subscription and a specialist teacher; policy pressure to show inclusion cheaply	Maphie (2025); Chirowamhangu (2024); Motitswe (2025); Sharma & Ha (2025)
P7	Learner agency is preserved when the learner controls the level of assistance, understands what the tool has done, and the assistance is faded on a plan; it is eroded when assistance is permanent, opaque and effort-minimising	Cognitive offloading; supported versus independent performance; sycophancy–dependence pathway	Darvishi et al. (2024); Fan et al. (2025); Stadler et al. (2024); Bastani et al. (2025); Chen (2026); Fan & Chang (2026); Choudhary et al. (2026)
P8	Assessment policy that bans AI without an accommodation pathway re-excludes disabled	Construct validity; reasonable-adjustment law; integrity regimes written before, or in	Morgan (2026); Kasneci et al. (2023); Bastani et al. (2025)

	students; policy that permits AI without objective specification empties accommodation of meaning; only a declared-use regime tied to stated objectives preserves both inclusion and validity	hasty response to, generative AI	
P9	In African education systems, and in schools above all, the inclusion dividend of generative AI is conditional on devices, connectivity, electricity, teacher AI literacy and language coverage, and the students most in need are those least likely to meet the conditions	Infrastructure and language as prior constraints on any AI-mediated accommodation (extrapolated from higher-education and non-African adoption studies)	Sharma & Ha (2025); Sevhake & Hofisi (2026); Kuyini (2026); Ng et al. (2023)

P1 to P3 classify what the tool is doing; P4 and P5 describe what the classification does to the people in the room; P6 lifts the logic to the school system; P7 and P8 state the conditions under which agency and validity survive; P9 states the conditions under which any of it reaches an African classroom.

4.2 DISCUSSION

The first thing to say about the propositions is which of them the evidence supports and which it merely permits. P7 is the best supported: Darvishi et al. (2024), Fan et al. (2025), Stadler et al. (2024), Bastani et al. (2025), Chen (2026) and Fan and Chang (2026) are, between them, a consistent body of experimental evidence that unstructured assistance raises supported performance while lowering, or failing to improve, independent performance and depth of reasoning, and that structure (guardrails, compulsory reflection, a division of labour that keeps evaluation with the student) narrows the gap, and Choudhary et al. (2026) supply a cross-sectional pathway by which an agreeable tutor is associated with dependence. None of them reports students with disabilities as a distinct group. P3 is therefore an inference, not a finding: it rests on the entanglement argument and on indirect evidence, from Alkhateeb (2026) that adaptive reliance is unevenly distributed among people with dyslexia and from the question Flütsch Keravec et al. (2026) put to speech-to-text. Seung and Basham (2026) pose the question for learning disabilities without yet answering it. The framework predicts that if the general experiments were replicated with disabled students the divergence would be larger, and that the tutoring guardrails that closed the gap in Bastani et al. (2025) would need redesigning so that the guardrail does not itself become a barrier. That is the most urgent testable prediction.

P1 is permitted by the evidence rather than established by it. Bray et al. (2024) found technology in UDL implementations used mostly to vary representation, with a clear gap on engagement and on action and expression; Hyatt and Owenz (2026) report that students in classes with a high proportion of disability chose evenly among UDL options for demonstrating learning in an AI-based assignment, and McCree et al. (2026) give special educators a framework for planning AI-supported instruction under UDL. What none of them measures is the quantity P1 turns on, which is whether the objective stayed fixed while the representation varied. A re-levelled text that preserves the ideas a unit is meant to teach and one that removes them look alike on the screen, and no study in the review distinguished them; until one does, P1 is a warning drawn from inclusive pedagogy's own history rather than a finding about the tools.

P2 is where the framework departs most from current practice. Accommodation decisions in most systems are made per student, recorded in a plan and applied across subjects. This cannot work for generative AI, because the classification of an AI use depends on the objective of the particular task, which varies by subject, by unit and sometimes by lesson. Morgan (2026) reached the same point from the law: a language model can only be classed as a reasonable adjustment once the institution has distinguished the competence standard it assesses from the method used to assess it, which is a statement about the task, not the student. Fan and Chang (2026) found the other half of the problem: even when the division of labour is prescribed, students with a high tendency to offload re-outsource the part they were assigned. A rule at the level of the plan is not enough; the objective has to be stated at the level of the task, and the student has to understand it, which brings the argument to P4.

P4 asks more of teachers than it might appear. Producing adapted material was for decades the visible work of the special educator; the early literature anticipated that AI would change that work (Marino et al., 2023), and later studies show the tools taking it on, from individualised education plan goals (Rakap & Balikci, 2026) to scaffolded instructional resources (van Garderen et al., 2025). The framework draws the opposite conclusion about the special educator's future: the scarce skill is now the ability to say what a task is for, which barriers to it are irrelevant and which are the point, and how much assistance a student should have this week and how much less next. Zraydi et al. (2026) found that teachers' intentions to use AI with students with learning disabilities were related to their AI-specific technological pedagogical content knowledge, and that only one component of that knowledge predicted them. Teacher-education systems that respond to generative AI by teaching prompt-writing will have trained for the wrong scarcity.

P5 is the proposition the technology literature finds hardest to see, because it is not about learning. Self-determination theory puts relatedness beside competence and autonomy as a need, and the synthesis by Fan et al. (2026) found that controlling or weakly structured environments may promote dependent cognitive offloading and over-reliance; the framework reads that finding socially as well as cognitively. A tool that lets a deaf student follow a lecture through captions, or a student with dyslexia read a text without a peer reading it aloud, removes a dependence that was often humiliating, and that is a gain for dignity. The same tool removes an interaction. Jolly et al. (2024) found that many deaf students eventually experienced effective communication access through real-time captioning but could initially struggle to overcome barriers to using it, and that the stories they told fell into four types, among them stories of resignation and stories of personal connection. A classroom in which every disabled student's support is a private device has achieved presence and lost something it does not yet know how to measure.

P6 is a prediction about policy rather than classrooms, and the paper's Africa framing rests on it. The inclusive-education literature from Tanzania, South Africa and Ghana (Maphie, 2025; Chirowamhangu, 2024; Motitswe, 2025; Malahlela & Johnson, 2024; Kuyini, 2026) describes one configuration: policy commitments to inclusion, placement of disabled learners in mainstream classes, and support that is thin because specialist teachers, devices, accessible buildings and small classes cost money the systems do not have. Into that configuration generative AI arrives as the first support that is cheaper than the problem; a ministry that cannot hire specialist teachers can license a chatbot. The framework predicts that AI accommodation will be adopted in such systems as a substitute for human support and physical accessibility, that the substitution will be presented as inclusion, and that the consequences for participation (P5) and independent capacity (P7) will be invisible in the outcome measures the systems use. The evidence is, at present, only the absence of any African study and Sharma and Ha's (2025) finding that country, social persuasion, perceived usefulness and perceived ease of use, rather than individual characteristics, predict adoption.

P8 restates the integrity problem from the disabled student's side. Kasneci et al. (2023) saw early that teachers would find it increasingly difficult to tell machine-generated text from student work, and that some institutions were already responding with bans; what they did not foresee was that integrity rules written in response would collide with disability law. A blanket prohibition on AI in assessed work is, for a student whose reasonable adjustment is an AI reader or scribe, a prohibition on the adjustment, and if Morgan's (2026) reading of UK equality law is right, a rule that forbids such an adjustment without distinguishing the competence assessed from the method would be hard to defend where disability rights are enforceable. Permitting AI generally dissolves the distinction between accommodation and substitution and with it any reason to record adjustments at all; Bastani et al. (2025) showed what unrestricted permission costs when the assessment is later taken without the tool. The framework's answer, a declared-use regime in which the objective of each assessed task is stated, the student declares which AI functions were used, and the use is adjudicated against the stated objective, is more demanding than either extreme. It is also the only policy consistent with both inclusion and validity, and examination bodies in African systems, which run some of the most prohibition-based integrity regimes anywhere, will have to confront it soon.

P9 is a reminder that none of this reaches a classroom without a device, a connection, electricity and a language the model handles well. Ng et al. (2023) set out the AI competencies teachers need and noted how few existing frameworks specify them; Sevhake and Hofisi (2026) found South African and Kenyan universities weighing efficiency against equity with

the infrastructure to do so unevenly distributed. The language point deserves more weight than the literature gives it. Generative models perform best in the languages on which they were trained most, and a student taught in Hausa, Yoruba, Igbo or Kiswahili will find the accommodation functions of the tool markedly weaker than a student in an English-medium school. The inclusion dividend will therefore go first to the urban, the English-medium and the connected, a distribution the inclusive-education movement will recognise, because it is the one it has been fighting.

Where the evidence contradicts the framework it is worth saying so. The meta-analytic finding that ChatGPT-based learning improves engagement (Xia et al., 2025) sits uneasily with a framework that warns of substitution: a tool that makes a disabled student want to try is not nothing. The reconciliation offered here is that engagement with the supported task and competence on the independent task are different outcomes, that Chen (2026) and Bastani et al. (2025) measured both and found them diverging, and that inclusive education should want both. The positive results for dyslexia (Gharaibeh & Basulayyim, 2025) and for visually impaired students (Elshaer et al., 2025) show that the accommodation and personalisation functions work; they say nothing yet about independent capacity over time, which is where the framework places the risk.

5. CONCLUSION AND RECOMMENDATION

Generative AI gives inclusive education what it asked for and then asks it a question it has not had to answer before: what, precisely, is the task? A reader, a scribe and extra time never forced that question because they could not do the task. A language model can, so every accommodation decision now contains an implicit statement about what the student is supposed to be learning, whether or not anyone has written it down. The framework makes that statement explicit. Personalisation that varies the route to a common objective, and accommodation that removes an irrelevant barrier, advance inclusion; substitution that performs the construct-relevant work undermines it; and which of the three a given use amounts to is a fact about the objective, not the tool.

Several recommendations follow. Ministries of education and curriculum agencies should require that each unit of a national curriculum state, for its principal tasks, which skills are the target and which are access skills; without that statement no teacher can classify an AI use and no examination body can adjudicate one. Examination councils, including the regional and national bodies of African systems, should replace blanket prohibitions on AI in assessed work with declared-use regimes tied to stated task objectives, and publish the accommodation pathway for students whose reasonable adjustment is an AI function before the first legal challenge rather than after it. Teacher-education institutions should move the preparation of general and special educators away from tool operation and towards objective specification, the design and fading of AI-supported accommodation, and the monitoring of independent alongside supported performance. School leaders should treat generative AI as a complement to, never a budget line replacing, specialist teachers, assistive devices and accessible buildings, and should review the social participation of AI-supported students and not only their outputs. Disabled people's organisations and parents' groups should be party to each of these decisions and should press the language question in particular, since an AI-in-education policy that ignores the tools' performance in the languages of instruction is a policy for the English-medium minority. Researchers should replicate the experiments on supported versus independent performance with disabled students as a distinct group, measure participation and belonging alongside learning, and run those studies in African schools, where the propositions of this paper will be tested by policy whether or not they are tested by research.

6. CONTRIBUTION TO EXISTING KNOWLEDGE

The paper's first contribution is conceptual. The literatures on generative AI in education and on inclusive education have proceeded without a shared vocabulary for deciding when AI assistance includes a disabled learner and when it merely produces an output for them. The personalisation–accommodation–substitution distinction supplies one, and the claim that the boundary between accommodation and substitution is drawn by the learning objective rather than the tool moves a question both literatures have been asking in the wrong place; it also explains why per-student adjustment plans cannot classify AI uses on their own.

The second contribution is a set of nine propositions specific enough to be tested and, in several cases, to be wrong. The predictions that the divergence between supported and independent performance will be larger for disabled students than for others,

that AI accommodation will be adopted in low-resource systems as a substitute for human support rather than a complement to it, and that privatised AI support reduces the peer interaction through which belonging is built are each falsifiable with designs that already exist in the general literature. The paper says where the current evidence supports, permits or contradicts each proposition, and treats the largest gap, the absence of any controlled study of generative AI with disabled learners in an African school, as a finding in its own right.

The third contribution is to policy in school systems where specialist support is scarcest. By lifting the accommodation–substitution distinction from the classroom to the system, the paper gives ministries, examination bodies and teacher-education institutions a way to recognise system-level substitution when it is presented as inclusion, and a set of requirements against which national AI-in-education policies can be judged. It offers inclusive-education practice in Africa a framework that is realistic about its constraints without conceding that cheapness is the same thing as inclusion.

7. REFERENCES

Alkahtani, B. N. (2024). The impact of artificial intelligence on quality of life for deaf and hard of hearing students. *American Annals of the Deaf*, *169*(4), 329-347. https://doi.org/10.1353/aad.2024.a946587

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*, 1-18. https://doi.org/10.1080/17483107.2026.2715246

Amal Krishnan, U. C., Lathabai, H. H., Kanjilal, R., Raman, R., & Nedungadi, P. (2026). Generative AI in disability-inclusive learning: a bibliometric and systematic literature analysis. *Disability and Rehabilitation: Assistive Technology*. Advance online publication. https://doi.org/10.1080/17483107.2026.2653073

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

Bray, A., Devitt, A., Banks, J., Sanchez Fuentes, S., Sandoval, M., Riviou, K., Byrne, D., Flood, M., Reale, J., & Terrenzio, S. (2024). What next for Universal Design for Learning? A systematic literature review of technology in UDL implementations at second level. *British Journal of Educational Technology*, *55*(1), 113-138. https://doi.org/10.1111/bjet.13328

Chen, X. (2026). Layer-sensitive cognitive offloading in generative AI-assisted writing: supported performance and independent no-AI outcomes. *Frontiers in Psychology*, *17*, 1906199. https://doi.org/10.3389/fpsyg.2026.1906199

Chirowamhangu, R. (2024). Inclusive education pandemic: Learning barriers for children with disabilities in South Africa. *African Journal of Disability*, *13*, a1462. https://doi.org/10.4102/ajod.v13i0.1462

Chiu, T. K. F., Xia, Q., Zhou, X., Chai, C. S., & Cheng, M. (2023). Systematic literature review on opportunities, challenges, and future research recommendations of artificial intelligence in education. *Computers and Education: Artificial Intelligence*, *4*, 100118. https://doi.org/10.1016/j.caeai.2022.100118

Choudhary, M. H., Dastgeer, G., Ahmed, S. A., Mohsin, M., & Naseem, S. (2026). How AI sycophancy shapes learner autonomy in digitalized learning: the mediating roles of metacognitive laziness and AI dependence. *Frontiers in Psychology*, *17*, 1957447. https://doi.org/10.3389/fpsyg.2026.1957447

Darvishi, A., Khosravi, H., Sadiq, S., Gašević, D., & Siemens, G. (2024). Impact of AI assistance on student agency. *Computers & Education*, *210*, 104967. https://doi.org/10.1016/j.compedu.2023.104967

Elshaer, I. A., AlNajdi, S. M., & Salem, M. A. (2025). Measuring the impact of large language models on academic success and quality of life among students with visual disability: An assistive technology perspective. *Bioengineering*, *12*(10), 1056. https://doi.org/10.3390/bioengineering12101056

Fan, M., & Chang, P. (2026). Telling students to evaluate does not make it happen: Task stage,

offloading tendency, and error detection in AI-assisted student writing. *Behavioral Sciences*, *16*(9), 1671. https://doi.org/10.3390/bs16091671

Fan, Y., Li, X., & Zhang, R. (2026). Family-school autonomy support for children's responsible use of generative artificial intelligence: a self-determination theory synthesis and developmental research agenda. *Frontiers in Psychology*, *17*, 1969217. https://doi.org/10.3389/fpsyg.2026.1969217

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, *56*(2), 489-530. https://doi.org/10.1111/bjet.13544

Flütsch Keravec, S., Schneider, H., & Hartmann, E. (2026). Writing acquisition and modality-effect of speech-to-text technology: An intervention study with students with dyslexia. *Reading & Writing Quarterly*, 1-19. https://doi.org/10.1080/10573569.2026.2664964

Gharaibeh, M., & Basulayyim, A. (2025). ChatGPT enhances reading comprehension for children with dyslexia in Arabic language. *Disability and Rehabilitation: Assistive Technology*, *20*(8), 2858-2871. https://doi.org/10.1080/17483107.2025.2508389

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

Hopcan, S., Polat, E., Ozturk, M. E., & Ozturk, L. (2023). Artificial intelligence in special education: a systematic review. *Interactive Learning Environments*, *31*(10), 7335-7353. https://doi.org/10.1080/10494820.2022.2067186

Hyatt, S. E., & Owenz, M. B. (2026). Using Universal Design for Learning and artificial intelligence to support students with disabilities. *College Teaching*, *74*(2), 89-96. https://doi.org/10.1080/87567555.2024.2313468

Jolly, A. J., Macfarlane, C. E., & Barker, B. A. (2024). Deaf/hard of hearing students' experiences with higher education's real-time captioning services. *Journal of Deaf Studies and Deaf Education*, *29*(3), 424-432. https://doi.org/10.1093/jdsade/enae019

Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, *103*, 102274. https://doi.org/10.1016/j.lindif.2023.102274

Kuyini, A. B. (2026). School principals' expectations of teachers to implement inclusive education practices and teachers' understanding of the expectations in Ghana. *International Journal of Inclusive Education*, *30*(11), 3306-3322. https://doi.org/10.1080/13603116.2025.2586073

Malahlela, M. K., & Johnson, E. (2024). South African teachers' application of inclusive education policies and their impact on learners with learning disabilities: Implications for teacher education. *Education Sciences*, *14*(7), 743. https://doi.org/10.3390/educsci14070743

Maphie, E. I. (2025). Implementation of inclusive education in Tanzania secondary schools: placements and supports for students with disabilities. *International Journal of Inclusive Education*, *29*(13), 2340-2357. https://doi.org/10.1080/13603116.2024.2348494

Marino, M. T., Vasquez, E., Dieker, L., Basham, J., & Blackorby, J. (2023). The future of artificial intelligence in special education technology. *Journal of Special Education Technology*, *38*(3), 404-416. https://doi.org/10.1177/01626434231165977

McCree, N., Buranova, N., Day, J., Juergensen, R., & Goran, L. (2026). Artificial intelligence in special education: Advancing inclusion through Universal Design for Learning, and culturally responsive teaching. *Journal of*

Special Education Technology. Advance online publication, Article 01626434261448827. https://doi.org/10.1177/01626434261448827

Morgan, D. (2026). Reasonable adjustments for neurodivergent students in higher education: generative AI, accessibility, and mediated authorship. *Disability & Society*, *41*(7), 1836-1856. https://doi.org/10.1080/09687599.2026.2667528

Motitswe, J. (2025). Teachers' perceptions on including learners with barriers to learning in South African inclusive education system. *African Journal of Disability*, *14*, a1543. https://doi.org/10.4102/ajod.v14i0.1543

Ng, D. T. K., Leung, J. K. L., Su, J., Ng, R. C. W., & Chu, S. K. W. (2023). Teachers' AI digital competencies and twenty-first century skills in the post-pandemic world. *Educational Technology Research and Development*, *71*(1), 137-161. https://doi.org/10.1007/s11423-023-10203-6

Rakap, S., & Balikci, S. (2026). Enhancing IEP goal development for preschoolers with autism: A preliminary study on ChatGPT integration. *Journal of Autism and Developmental Disorders*, *56*(4), 1682-1687. https://doi.org/10.1007/s10803-024-06343-0

Ronksley-Pavia, M., Nguyen, L., Wheeley, E., Rose, J., Neumann, M. M., Bigum, C., & Neumann, D. L. (2025). A scoping literature review of generative artificial intelligence for supporting neurodivergent school students. *Computers and Education: Artificial Intelligence*, *9*, 100437. https://doi.org/10.1016/j.caeai.2025.100437

Seung, Y., & Basham, J. D. (2026). Cognitive offloading in the age of generative AI: What does it mean for students with learning disabilities? *Learning Disability Quarterly*, *49*(3), 121-133. https://doi.org/10.1177/07319487261439132

Sevhake, M. G., & Hofisi, C. (2026). Navigating AI in higher education: Balancing efficiency, equity, and autonomy in South Africa and Kenya. *AI in Education*, *2*(3), 24. https://doi.org/10.3390/aieduc2030024

Sharma, S., & Ha, L. (2025). Generative AI in higher education: A comparative study of ChatGPT adoption, perception, and use among college students in the Global North and Global South. *Human-Machine Communication*, *11*, 107-126. https://doi.org/10.30658/hmc.11.6

Stadler, M., Bannert, M., & Sailer, M. (2024). Cognitive ease at a cost: LLMs reduce mental effort but compromise depth in student scientific inquiry. *Computers in Human Behavior*, *160*, 108386. https://doi.org/10.1016/j.chb.2024.108386

Urbina, J. T., Vu, P. D., & Nguyen, M. V. (2025). Disability ethics and education in the age of artificial intelligence: Identifying ability bias in ChatGPT and Gemini. *Archives of Physical Medicine and Rehabilitation*, *106*(1), 14-19. https://doi.org/10.1016/j.apmr.2024.08.014

van Garderen, D., Lannin, J., Rodrigues, J., Smith, C. M., & Shahale, F. V. (2025). Using generative AI to create scaffolded instructional resources for students with disabilities. *Journal of Special Education Technology*, Advance online publication. https://doi.org/10.1177/01626434251408056

Xia, Q., Li, W., Yang, Y., Weng, X., & Chiu, T. K. F. (2025). A systematic review and meta-analysis of the effectiveness of Generative Artificial Intelligence (GenAI) on students' motivation and engagement. *Computers and Education: Artificial Intelligence*, *9*, 100455. https://doi.org/10.1016/j.caeai.2025.100455

Zhai, C., Wibowo, S., & Li, L. D. (2024). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: a systematic review. *Smart Learning Environments*, *11*(1), 28. https://doi.org/10.1186/s40561-024-00316-7

Zraydi, F., Opoku, M. P., Guirguis, B. M., Ndijuye, L. G., & Gyimah, E. M. (2026). Generative AI intentions and technological pedagogical and content knowledge: teachers' readiness to use AI for students with learning disabilities and difficulties. *Disability and Rehabilitation: Assistive Technology*, 1-17. https://doi.org/10.1080/17483107.2026.2714681