Routledge
Taylor & Francis Group

International Journal of Inclusive Education

ISSN: 1360-3116 (Print) 1464-5173 (Online) Journal homepage: www.tandfonline.com/journals/tied20

AI in one hand, inclusive education in the other: why mainstream teachers who use AI do not use it for inclusive education

Roberta Besozzi & Anoop Menon

INTERNATIONAL JOURNAL OF INCLUSIVE EDUCATION
https://doi.org/10.1080/13603116.2026.2715640

AI in one hand, inclusive education in the other: why mainstream teachers who use AI do not use it for inclusive education

Roberta Besozzi [a] and Anoop Menon[b]

[a]Department of Teacher Education, Université de Fribourg, Fribourg, Switzerland; [b]Eisengard AI, San Francisco, CA, USA

ABSTRACT

Mainstream Swiss teachers have adopted generative AI rapidly, yet rarely turn to it when facing inclusion difficulties – even though the same tool that drafts their lesson plans could equally suggest strategies for a specific student. The gap between broad adoption of digital tools and their application to specific educational purposes is well documented. Still, existing explanations were developed for technologies whose domains and functions were pre-specified by design. Generative AI has no such pre-specified domains, raising an open question: why does the gap persist? This mixed-methods exploratory study addresses that question with a trilingual survey of 69 Swiss mainstream teachers and nine in-depth interviews with teachers and school leaders across three cantons. We find a steep behavioural cascade: 74% of teachers use AI professionally, 18% extend it to inclusion-specific purposes, and only 4% name it as a coping strategy for inclusion difficulties. Three factors are associated with crossing to inclusion-specific use: breadth of AI engagement, materials creation as a bridge to student-specific adaptation, and professional distress. We theorise the pattern as a cognitive category boundary and organise the findings through the Boundary Model of AI in Inclusive Education.

ARTICLE HISTORY
Received 27 April 2026
Accepted 2 August 2026

KEYWORDS
Generative AI; inclusive education; teachers; the boundary model of AI in inclusive education; mixed methods; Switzerland

1. Introduction

The gap between broad technology adoption by teachers and its application to specific educational purposes has been extensively documented in educational technology research (Cuban 2001; Davis 1989; Ertmer and Ottenbreit-Leftwich 2010). This pattern appears in inclusive education as well. Okolo and Diedrich's (2014) statewide survey of 1143 educators found that teachers use technology frequently in their personal and professional lives, yet substantially less during the instruction of students with special educational needs. Alvarez (2024) frames this phenomenon theoretically through the capability approach, arguing that digital tools do not automatically convert into inclusive capability.

CONTACT Roberta Besozzi roberta.besozzi@unifr.ch Department of Teacher Education, Université de Fribourg, Rue de Morat 36, Fribourg 1700, Switzerland

The candidate explanations for this gap come from three research traditions, which §2 develops. Educational technology theorises teachers' beliefs as a second-order barrier that persists after first-order conditions such as access and training are addressed (Ertmer 1999; Ertmer and Ottenbreit-Leftwich 2010). The information systems literature describes a parallel phenomenon of post-adoption narrowing, in which users of complex tools default to a small, stable set of features (Jasperson, Carter, and Zmud 2005). Goodhue and Thompson (1995) add a task-technology fit account, according to which adoption converts into sustained use only where users perceive a genuine fit between tool and task. These accounts are not mutually exclusive; a teacher's beliefs about what a technology is for shape whether they perceive a good fit with a given task. What they share is a common referent – each was developed for technologies with pre-specified domains.

Generative AI disrupts that common referent. It has no designed-in domain and no pre-specified function; the same prompt interface could draft a lesson plan, interpret a diagnosis, or generate strategies for a specific student. Whether the mechanisms above still apply, or whether new ones emerge, remains an open empirical question. We work through each framework in this light in the literature review. Here we simply note that the adoption-application question warrants renewed examination for a tool without built-in domains, and for the mainstream teachers who carry the everyday load of inclusive classrooms.

Empirical work on generative AI and inclusive education is recent but converging. Weekly AI use among special education teachers in Wisconsin sits below 15% (Naatz and Ruppar 2025); Chinese special education teachers rate generative AI as less useful for individual education plans than for lesson planning (Lin, Cui, and Zhu 2026); across 55 economies, support for students with special educational needs remains a minority application of AI in teaching, and teachers' AI beliefs emerge as a significant mediator of the effect of professional development on AI integration (Liu et al. 2026; OECD 2026). These studies focus on specialists and report adoption frequencies; none examines the within-teacher pathway from general AI use to inclusion-specific application among mainstream teachers.

This article addresses that gap. Among mainstream teachers who have already adopted AI in their professional work, under what conditions does this adoption extend to inclusion-specific application and when it does not, what accounts for the gap? Throughout, we distinguish general professional AI use (lesson planning, materials creation, administrative tasks) from inclusion-specific use (understanding a diagnosis or difficulty, or finding strategies for a specific student); §3.5 details the operationalisation. Four sub-questions structure the analysis: (i) who, among mainstream teachers, adopts AI? (ii) among those who adopt, who extends use to inclusion-specific purposes? (iii) what do teachers do when they face inclusion difficulties and where does AI enter that repertoire? (iv) how should this distribution be theorised?

To answer these questions, we use a convergent mixed-methods design. The quantitative core is a trilingual survey (French, Italian and German) of 69 Swiss mainstream teachers, collected in March and April 2026. The qualitative material comes from nine in-depth interviews with teachers and school leaders conducted in a parallel programme (Besozzi and Menon n.d.) and from free-text responses to the survey. The survey describes structural patterns, while the qualitative material provides interpretive access to the professional logics that give those patterns their meaning (§3.1).

The Swiss setting chosen for this mapping is itself consequential. Here, inclusive education is a live and contested field of teacher work, while generative AI use among teachers is rising rapidly (OECD 2026). Yet national policy remains fragmented across cantons, no framework addresses the AI-inclusion intersection (Bešić, Schulz, and Schmid-Meier 2025), and no AI tool currently exists for mainstream teachers supporting students with special educational needs.

From our survey, we find that of 69 mainstream teachers, 74% use AI professionally, 18% extend it to inclusion-specific purposes and only 4% name AI as a strategy when facing inclusion difficulties. We label these three steps adoption, extension and coping. No dispositional or structural variable shows a robust association with adoption. At extension, three factors are associated with crossing: breadth of AI engagement, materials creation and professional distress. At coping, teachers overwhelmingly name colleagues, specialists and documentation as their resources, rather than AI.

We theorise this pattern as a cognitive category boundary. Inclusive education and AI occupy separate compartments in teachers' professional thinking, a boundary that is cognitive, not technological, since the same generative AI tool can serve both purposes. The findings are organised through the Boundary Model of AI in Inclusive Education, which traces the three transitions from general AI use to inclusion-specific application.

The remainder of the article is structured as follows. Section 2 reviews the relevant literatures across three traditions – educational technology, the organisation of professional work and generative AI in inclusive education. Section 3 describes the mixed-methods design, the survey instrument, the sample and the analytical approach. Section 4 presents the results, opening with the Boundary Model of AI in Inclusive Education that structures the findings across three transitions. Section 5 discusses the implications of the cascade, the cognitive category boundary and what does and does not predict AI use, before situating the findings in an international context. Section 6 concludes.

2. Literature review

2.1. Mechanisms of the adoption-application gap

Ertmer's (1999) distinction between first- and second-order barriers established that cognitive and dispositional factors persist after access and training are addressed. In inclusive education specifically, Katz (2015) found that implementation of Universal Design for Learning (UDL) supported by structured professional development improved teacher self-efficacy in inclusive classrooms, and Capp's (2017) meta-analysis of UDL research points to similarly conditional effects. Beliefs, in the updated empirical account, operate as a lens or filter shaping what teachers perceive a technology as being for (Ertmer and Ottenbreit-Leftwich 2010). Teachers' prior beliefs about generative AI, formed in contexts where AI was positioned for lesson planning and content creation tools, may filter out its potential application to inclusion work.

Standard adoption models (Davis 1989; Venkatesh et al. 2003) predict whether users adopt a technology, but do not address what happens after adoption. Jasperson, Carter, and Zmud (2005) describe a parallel phenomenon in the information systems literature – post-adoption narrowing, in which users of complex tools default to a small and stable subset of features unless organisations intervene deliberately. They articulated this mechanism for tools with pre-specified feature sets, so it does not map directly onto generative AI. The underlying claim, that routinisation narrows engagement, may still apply, operating through conceived use case rather than feature selection.

Goodhue and Thompson (1995) offer a third account in the task-technology fit tradition. Technology affects professional performance only where users perceive a good fit between task and tool. For inclusion work (seen as relational, contextual and judgement-intensive) perceived fit with a general-purpose text generator may be low on rational grounds.

2.2. The organisation of the professional work of teachers

The sociology of teachers' work offers a complementary set of explanations that do not depend on cognition. Abbott (1988) shows that professions are structured by jurisdictions. Specialists claim particular domains of work, mainstream practitioners defer to that allocation. Applied to inclusive education, this predicts that mainstream teachers will orient to colleagues, specialists and specialist-provided resources when facing inclusion difficulties, rather than to tools that might encroach on specialist territory.

The division of labour within teaching tends to keep instruction and specialised support in distinct compartments (cf. Tardif and Lessard 1999, on the structure of teachers' work). Teachers mentally file their tools within one domain or the other. AI has been filed under lesson planning and content creation, and does not migrate spontaneously into the specialised-support domain even though it could. Clot's (2008) concept of activité empêchée captures a further dimension – the professional activity teachers would undertake but structurally cannot. In inclusive education, this activity is relational in character, reinforcing the reliance on relational rather than technical supports. Together with the cognitive mechanisms reviewed above, these accounts suggest that an adoption-application gap for inclusive education would be sustained at both individual and institutional levels.

2.3. Generative AI and inclusive education

Empirical work on generative AI and inclusive education has, as noted in the Introduction, concentrated on specialist populations. Naatz and Ruppar's (2025) exploratory survey of special education teachers in Wisconsin found that fewer than 15% use AI weekly, with frequency associated with self-efficacy and perceived usefulness. Lin and colleagues (2026) interviewed Chinese special education teachers and found that generative AI was rated lower for the preparation of individual education plans than for lesson planning, a finding that echoes the adoption-application gap for inclusion-specific tasks.

At population level, the OECD's (2026) Digital Education Outlook reports that support for students with special educational needs remains a minority application of AI among lower-secondary teachers across the 55 economies surveyed in TALIS 2024 (roughly a third of AI-using teachers). Liu and colleagues (2026), using TALIS 2024 microdata (n = 34,628), find that teachers' AI beliefs outperform professional development as a predictor of AI integration. Swiss-specific evidence on pre-service teachers (Alvarez et al. 2024) documents rapid generative AI acceptance, with ease-of-use a weaker predictor than for earlier technologies. Alvarez (2023) reminds us that inclusion itself is contested – integration (the learner adapts to an existing structure) must be distinguished from inclusion (the structure is reconfigured). Digital artefacts can operate on either side of this distinction. Weber et al. (2022) make this tension concrete, noting that digital developments simultaneously create opportunities for inclusion and risks of new exclusions. An emerging strand of work documents the assistive potential of generative AI for accessibility and equity in inclusive settings, to date centred on higher-education students with visual impairments (Khlaif et al. 2025). The within-teacher pathway from general AI adoption to inclusion-specific application among mainstream teachers remains unexamined.

3. Methods

3.1. Design

This study uses a convergent mixed-methods design (Creswell and Plano Clark 2018) in which a quantitative survey strand and a qualitative strand were collected independently and brought together at interpretation. The design follows from the research question: establishing whether an adoption-extension gap exists, and how wide it is, requires structured measurement across many teachers, while understanding why it persists requires access to how teachers themselves categorise AI and inclusion work. The survey strand is primary for the first task; the qualitative strand (in-depth interviews from a parallel programme and open-ended survey responses) serves an interpretive and illustrative role for the second. The strands were integrated at three points: the thematic structure of the qualitative programme informed the survey instrument (§3.2), qualitative excerpts were selected to illustrate and probe patterns established quantitatively (§3.7) and the Boundary Model itself was generated by reading the two strands against each other (§4.1). Figure 1 summarises the design.

3.2. Survey instrument and construction

The survey comprises 25 closed items and one optional open-ended item. Three sources shaped the instrument. First, the thematic structure emerging from the qualitative

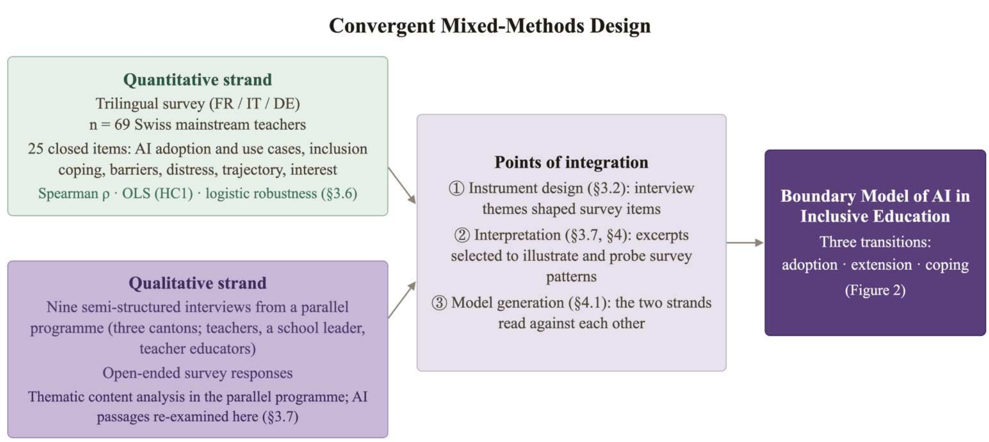

Figure 1. The convergent mixed-methods design: data sources, analyses, and points of integration.

programme associated with (Besozzi and Menon n.d.) identified five domains in teachers' accounts of inclusive education and digital tools. Second, Ertmer's (1999) barrier distinction and its operationalisations provided a conceptual framework. Third, items were adapted from validated instruments for self-efficacy in inclusive practices. The team prepared a trilingual version, with native-speaker review before distribution. The instrument was refined iteratively based on early respondent feedback.

3.3. Survey sample

For the survey, we recruited the sample through snowball sampling (Noy 2008), beginning with initial seeds in multiple Swiss cantons and distributing the survey through professional association channels. Sixty-nine teachers responded across three language regions – 42 from the French-speaking region, 25 from the Italian-speaking region and 2 from the German-speaking region. By school level, 35 teach in primary schools, 12 in lower secondary (Sec I) and 22 in upper secondary (Sec II). The sample consists of mainstream teachers responsible for inclusive classrooms; specialised teachers are not part of the sampling frame. All respondents hold a cantonally recognised teaching qualification. The two German-speaking respondents are retained in aggregate analyses and flagged throughout as insufficient for separate regional inference.

French- and Italian-speaking respondents are analysed jointly. The two regions share a common pedagogical tradition, and pilot analyses found no substantive differences in either the distribution of inclusive-education responses or the predictor structures we examine. Full regional stratification would require a substantially larger sample and lies beyond the scope of this study; we return to this point in the limitations.

3.4. Ethics and consent

The quantitative part of the study used a non-interventional, anonymous online survey. Swiss federal legislation on research involving human subjects (Federal Act on Research involving Human Beings, HRA, Art. 2) exempts anonymous surveys that neither collect health-related data nor pose any risk to participants from ethics committee review. Participants gave informed consent through an introductory screen before beginning the survey. We collected no personally identifiable data. For the qualitative part, written informed consent was obtained for interview data from the parallel programme.

3.5. Survey measures

AI adoption was measured on a four-point ordinal scale (never / once or twice / occasionally / regularly). AI use cases were captured by a multi-select item with six options: lesson planning, materials creation, administrative tasks, understanding a diagnosis or difficulty, finding strategies for a specific student and other. We defined inclusion-specific AI use as selecting at least one of the two inclusion-oriented options (understanding a diagnosis, or strategies for a specific student). The inclusion gradient codes these as a variable with four levels: no inclusion-specific use, informational only, applied only, or comprehensive (informational and applied).

Inclusion coping strategies were measured by an item asking, ‘When you face difficulties related to the inclusion of a student, what do you typically do?’ with six options (up to three selections), including ‘use an AI tool’. Barriers to AI use among non-users were captured by a multi-select item. Professional distress was measured by an item on prevented professional activity (five-point Likert), self-efficacy for inclusive practices (Bandura 1997) by a four-point Likert item, professional trajectory by a single item on the perceived evolution of professional ease since beginning work with students with diverse needs (five-point Likert: significantly decreased to significantly increased), and unmet need by a composite index of unmet support across challenge categories. Interest in an inclusion-specific AI tool was measured by a single item asking teachers how interested they would be in an AI tool designed to help teachers support students with diverse needs, on a four-point scale (not at all to very interested).

3.6. Survey analysis

The primary specification uses Spearman rank correlations for bivariate associations and OLS regression for multivariate models. All multivariate p-values reported in the results are based on heteroskedasticity-consistent (HC1) robust standard errors. Logistic regression serves as a robustness check for binary outcomes. School level and language region are included as controls in multivariate specifications. All associations are cross-sectional; no causal inference is warranted.

With 51 AI users, the study provides 57% power to detect a medium effect ($\rho? = 0.30$) and 80% power at or above $\rho? = 0.38$, at $\alpha = 0.05$ (two-tailed). Null findings below $\rho? = 0.38$ should therefore be read as consistent with a range of underlying effect sizes, not as strong evidence of no effect. Cascade proportions are more precise: Wilson 95% CIs are 63–83% for adoption, 10–30% for inclusion-specific use, and 2–12% for AI as inclusion coping.

3.7. Qualitative material

Qualitative material is drawn from two sources, each serving a specific interpretive role. First, nine semi-structured interviews were conducted in a parallel programme (Besozzi and Menon n.d.) with teachers, a school leader and teacher educators across three cantons, spanning early-childhood to upper-secondary and tertiary teaching. The interview guide comprised seven open questions, including the sense of professional competence in the face of inclusion, available tools and training, institutional support systems and the use of artificial intelligence and its limits. Interviews were analysed through thematic content analysis in the parallel programme (Miles, Huberman, and Saldaña 2020); for the present article, the passages concerning digital tools and AI were re-examined in light of the survey patterns, and the excerpts quoted in §4 were drawn from this material. Second, open-ended responses in the present survey offer first-person reflection on distress, coping, and the limits of available support. The excerpts quoted in §4 are illustrative rather than representative; they were chosen as the clearest articulations of stances observed in the wider material and, in the case of Marina and Camille (§4.4), as contrast cases marking the two sides of the cognitive category boundary. All interviewees are referred to by pseudonyms; where the same participant appears in the parallel paper, pseudonyms are aligned. Written informed consent for the interview material is described in §3.4.

4. Results

We now introduce the Boundary Model of AI in Inclusive Education and then trace its three transitions empirically: adoption (who uses AI), extension (who applies it to inclusion), and coping (what teachers do when facing inclusion difficulties).

4.1. Conceptual framework: boundary model of AI in inclusive education

The Boundary Model of AI in Inclusive Education (Figure 2) organises the findings into three transitions, drawing its interpretive vocabulary from the mechanisms reviewed in §2. The first concerns adoption; we find that dispositional signals (self-efficacy, experience) are weak at best, and structural variables show no association with AI use. The second concerns extension to inclusion-specific use, where we see that only teachers who use AI for materials creation reach it. The third concerns how far they go once there, where breadth of AI engagement and professional distress both play a role. At each transition, measured need, workload and structural conditions are not associated with movement (we return to this in §5.3). Arrows in Figure 2 indicate statistical association rather than causation. We note that the model was generated from the present findings rather than derived from prior theory or tested against new data; it is offered as a tentative analytical device that future research can test, refine, or reject.

In the vocabulary of the capability approach, AI adoption is the capability and inclusion-specific use is the functioning. Three factors appear to convert one into the other – breadth of AI engagement, which provides structural exposure to a wider range of possible uses, materials creation, which sits cognitively close to student-specific adaptation and professional distress, which provides the motivational push to look for new solutions.

With generative AI, the domain boundary ceases to be a property of the tool and becomes a property of the encounter between the teacher and the tool. The question is not which features of the tool are used but which of its potential uses have entered the professional imagination of the teachers. We name this a cognitive category boundary; inclusion and AI occupy separate compartments in professional thinking, not because the technology enforces separation but because the conditions under which teachers form beliefs about what AI is for have not generated an integrated category.

Figure 3 summarises the sharp falloff, the behavioural cascade, in our data at each of these three transitions.

4.2. Adoption: who uses AI and who does not

4.2.1. Adoption is decoupled from inclusion conditions

Two patterns emerge at the adoption boundary, neither strong. Dispositional signals point in plausible directions but do not reach statistical significance consistently. Self-efficacy for inclusive practices is positively associated with AI use across tests but with

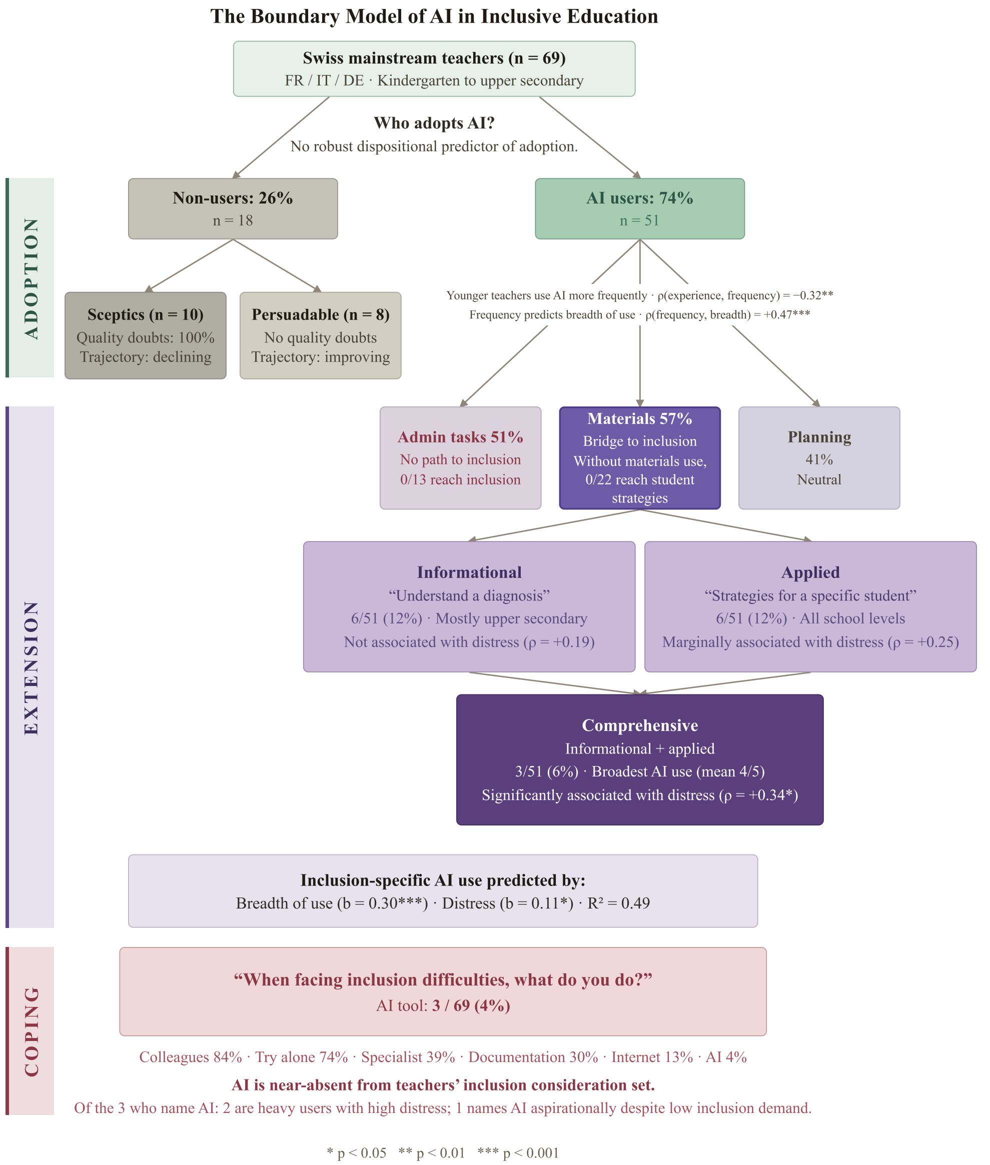

Figure 2. The boundary model of AI in inclusive education. Three transitions are indicated.

mixed strength (Spearman $\rho? = +0.22$, $p = 0.07$; logistic regression OR = 2.55, $p = 0.046$). Experience is negatively associated with frequency of use ($\rho? = -0.32$, $p = 0.009$), and primary-level teaching with lower AI frequency ($b = -0.54$, $p = 0.055$), though primary teachers in our sample are also more experienced on average. A joint OLS model (AI frequency = self-efficacy + experience), controlling for school level and language region, yields $R^2 = 0.15$, with experience marginal ($p = 0.084$) and self-efficacy not significant ($p = 0.171$); the dispositional signals weaken once structural variables are accounted for.

Structural variables, by contrast, show no association with AI engagement. Unmet inclusion need ($\rho? = +0.09$), prevented activity ($\rho? = -0.09$), hours on inclusion

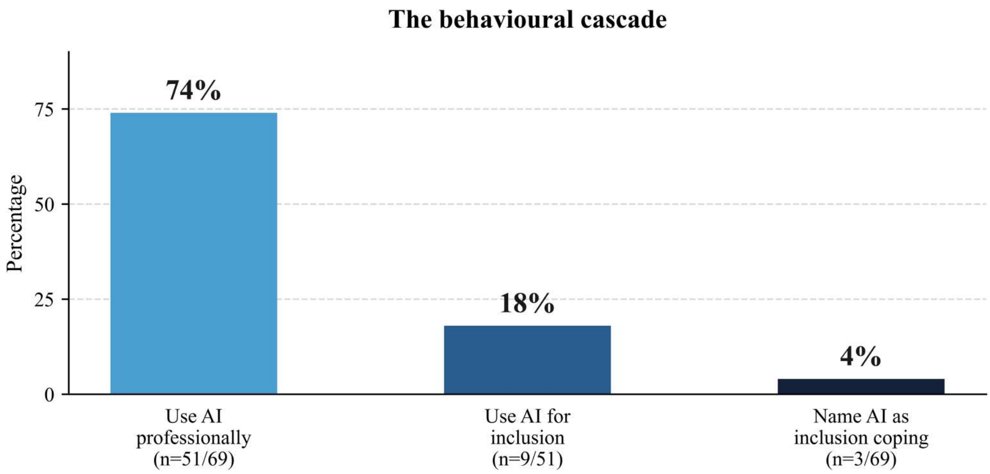

Figure 3. The behavioural cascade. Of 69 mainstream teachers, 51 use AI professionally (74%); of those, 9 apply AI to inclusion-specific purposes (18%); and only 3 of the 69 (4%) select AI when asked about coping strategies for inclusion difficulties.

outside contracted time ($\rho? \approx 0$), workplace flexibility ($\rho? = +0.15$), professional trajectory ($\rho? \approx 0$), specialist access, class size and training variables show no association with AI engagement (all $p > 0.13$). Teachers in large classes with heavy inclusion burdens do not use more or less AI than teachers with lighter loads. These patterns hold across alternative codings of the AI frequency variable.

This pattern can be interpreted in two ways (which we develop in Section 5.3). Disposition may simply not matter much. Or, structure shapes the cognitive frame in which teachers form a sense of what AI is, prior to and distinct from the moment of adoption; disposition, in this interpretation, is co-evolving with structure, not its alternative.

4.2.2. Non-users: quality concerns and exhausted trust

Among the 18 teachers who do not use AI, barriers are predominantly second-order. 61% cite belief-based barriers (quality doubts, privacy concerns), and 44% cite first-order barriers (time, lack of knowledge). Quality doubts track declining professional trajectories, where six of eight non-users with declining trajectories cite quality doubts, compared with one of six with improving trajectories ($\rho? = -0.47$, $p = 0.05$). Two subgroups emerge – sceptics ($n = 10$) report quality doubts, declining trajectories and low interest in AI tools; the persuadable ($n = 8$) cite time constraints rather than quality concerns, report improving trajectories and would be most accessible to targeted professional development. A primary teacher with twenty-five years of experience, in the sceptical subgroup, articulates a mood that extends beyond AI:

‘Magic wands (all proposed solutions, including those from training) are disconnected from reality, they are nothing but band-aids … My doctor “pulled me out of the deep fryer”?, that says it all!’ – Survey respondent, primary teacher, 25 years’ experience, French-speaking Switzerland.

The distrust expressed here targets not AI specifically but the entire class of proposed solutions. Read against Ertmer's second-order barrier, it reflects a belief formed across years of unmet promises rather than a position on any particular tool.

4.3. Extension: who applies AI to inclusion-specific purposes

4.3.1. The use-case fork

Among the 51 AI users, the most common use cases are materials creation (57%), administrative tasks (51%) and lesson planning (41%). Inclusion-specific options are uncommon (12% each: understanding a diagnosis, finding strategies for a specific student). The base rates matter less than the conditional structure. Among the 15 teachers who use AI for planning and/or administration but not for materials, none use AI for inclusion. The pattern extends to student-specific strategies; among the 22 teachers who do not use AI for materials, none use AI for student strategies. A single teacher in this group uses AI only to understand diagnoses. Every use-case combination that includes materials has a non-zero inclusion-specific rate. Inclusion rates also rise with the number of general use cases selected – 11% for one use case, 15% for two, 50% for three (planning, materials, administration).

These are associations, not causal claims. Materials use and inclusion-specific use may both reflect a content-focused orientation rather than one causing the other. But the conditional probabilities are empirical facts, consistent with Ertmer and Ottenbreit-Leftwich's (2010) argument that content-specific engagement enables domain transfer whereas administrative use does not.

4.3.2. Two orientations: informational and applied

The nine teachers who use AI for inclusion-specific purposes divide into two empirically distinct groups, informational and applied, with a comprehensive subgroup that spans both.

Informational use refers to teachers who select 'understanding a diagnosis or difficulty' ($n = 6$). Five of six are upper-secondary teachers. Informational use shows no association with distress ($\rho? = +0.19$, $p = 0.19$).

Applied use refers to teachers who select 'strategies for a specific student' ($n = 6$). Applied use is distributed across school levels, including the three primary teachers in our inclusion-specific subgroup and shows a marginal association with distress ($\rho? = +0.25$, $p = 0.087$).

Three teachers select both, forming a comprehensive subgroup ($n = 3$). They span early to mid-career (mean experience 12 years), report broad AI engagement (mean 4 of 5 use cases) and report high distress (mean 4.3 on the five-point measure; $\rho? = +0.34$, $p = 0.018$), the strongest distress association among the three subgroups.

4.3.3. The inclusion gradient

We code these distinctions as an ordinal gradient capturing depth of boundary-crossing: general use only ($n = 42$), single-pathway inclusion use (informational or applied, $n = 6$) and comprehensive use ($n = 3$). The two single-pathway groups are qualitatively distinct (§4.3.2) but represent equivalent steps across the boundary; the gradient orders by depth of crossing, not by orientation. Breadth of AI use is associated with movement along it (b

= +0.30, $p < 0.0001$), as is distress ($b = +0.11$, $p = 0.04$), controlling for school level and language region ($R^2 = 0.49$). Each additional general use case corresponds to a 0.30-point shift along the three-step gradient. With only six single-pathway and three comprehensive users, these estimates characterise the gradient's shape rather than measure precise effects. Figure 4 illustrates these patterns.

The variables that organise teachers' inclusion experience, such as unmet need, structural flexibility, workload, professional trajectory, school level, language region, show no association with the gradient, just as they showed none at adoption. Inclusion experience and AI engagement remain disconnected, a pattern we take up in §5.3.

4.4. Coping: what teachers do when facing inclusion difficulties

Asked what they typically do when facing inclusion difficulties, teachers select colleagues (84%), trying on their own (74%), contacting specialists (39%), consulting documentation (30%) and the internet (13%). 'Use an AI tool' is selected by only three of the 69 respondents (4%). Up to three options were permitted, so this is not a ceiling effect. The remaining 48 AI users do not name AI when asked about coping. Among the 52 teachers who express interest in an inclusion-specific AI tool, only three name AI as a current strategy. The coping repertoire is dominated by relational resources; §4.4.2 profiles the three teachers who do name AI.

When most teachers think 'inclusion difficulty', the evaluation step ('could AI help?') does not occur. The category 'AI as lesson planning' filters out 'AI as inclusion support' before an evaluation is made. Two interview extracts, selected as contrast cases marking the two sides of the boundary (§3.7), specify what this filtering looks like in professional terms.

'If I have to do something, I need to think about my students, and it [AI] doesn't know my students'. – Marina (interviewee), early primary teacher, Italian-speaking Switzerland.

Marina contrasts AI, which she describes as generic and 'without emotions' with a specialist who observed her classroom and provided specific, actionable feedback. For

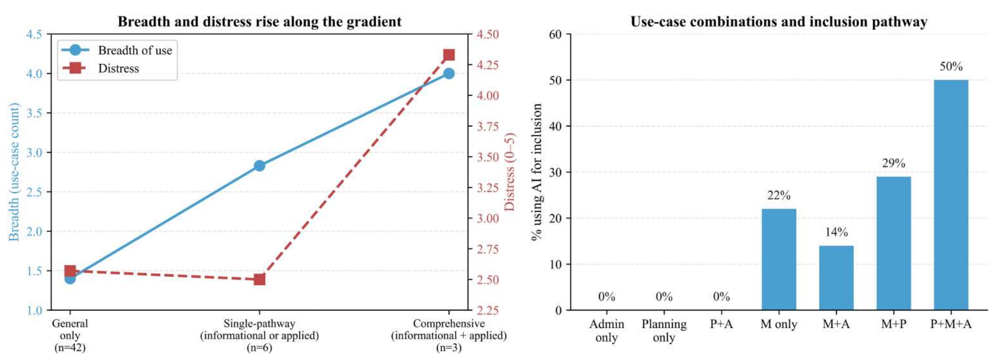

Figure 4. Left: breadth of AI use and distress rise across the three-step inclusion gradient (general / single-pathway / comprehensive); see §4.3.2 for the qualitative distinction between informational and applied within single-pathway use. Right: conditional probability of inclusion-specific AI use by general use-case combination. No teacher who uses AI without materials reaches student-specific use; a single teacher who uses AI without materials reaches informational use.

her, inclusion work is irreducibly particular, while AI is irreducibly general. This is not provisional scepticism but a professional conviction about what kind of knowledge inclusion work requires.

'I would probably ask AI for suggestions to personalise learning paths based on students' specific needs and abilities'. – Camille (interviewee), primary teacher, Italian-speaking Switzerland.

Camille already uses AI for materials creation and can describe a plausible next step, using AI for student-specific personalisation, but has not taken it. Luca, the school leader among our interviewees, confirmed the pattern from an institutional vantage point; he was not aware of any teachers in his schools who used AI for inclusion. Between Marina's conviction and Camille's articulated-but-untaken next step lies the cognitive category boundary this paper has sought to describe.

4.4.1. Interest without application

Seventy-five percent of teachers in the sample express interest in an inclusion-specific AI tool. Yet expressed interest is associated with current general AI use ($\rho ? = +0.52$, $p < 0.001$), not with unmet need ($p = 0.31$) or distress ($p = 0.50$). Teachers express interest not because they experience the absence of the tool, but because they already use AI elsewhere. The desire for an inclusion-specific tool exists, the mental category in which the existing tool could play this role does not.

4.4.2. Profiles of those who cross the boundary and those who do not

The nine teachers who use AI for inclusion are, on average, of mid-career experience (mean 12.9 years' experience), more frequent AI users, broader in their AI repertoire (3.22 use cases on average), and report moderate self-efficacy (2.89 out of 4). Eight of nine have no formal special educational needs training. Two subgroups emerge. High-distress, high-agency teachers (between four and eight years of experience, regular AI users, four to five use cases, frequent prevented activity) deploy AI alongside other resources to address an acute problem. A smaller group of lower-distress explorers approach informational use from a distance.

The contrast group sits on the other side. Twenty AI users report weekly prevented activity in inclusion and do not use AI for inclusion at all. Fourteen of the twenty are primary teachers; their AI use is less frequent and narrower in repertoire. For these teachers, the cognitive category boundary exacts its highest cost – they have AI in their hands yet do not apply it to what causes them the greatest professional difficulty. What separates them from the nine is not distress, need, or training, but breadth of engagement.

Three teachers select AI when asked what they do facing inclusion difficulties. Two also use AI for inclusion-specific tasks (one comprehensive, one applied), the deepest engagement in the sample, where accumulated breadth and distress have brought AI inside the inclusion repertoire. The third is a primary teacher with minimal inclusion difficulty (low distress, no unmet need, high specialist access) who has not applied AI to inclusion-specific tasks but names it aspirationally. All three report improving professional trajectories, the only subgroup uniformly on the improving end of the scale. The boundary can therefore be crossed by accumulated engagement under distress, or by anticipatory cognitive openness in its absence.

5. Discussion

5.1. Reading the cascade through the framework

The behavioural cascade (74% adoption, 18% extension, 4% coping) is what we observe, while the Boundary Model of AI in Inclusive Education explains why it takes this shape, naming the three transitions – adoption, extension, coping – at which the cognitive category boundary either holds or is crossed. Adoption itself is not robustly predicted by any single variable, a point we return to in §5.3. The interesting action emerges at the subsequent transitions. The use-case fork echoes Ertmer and Ottenbreit-Leftwich's (2010) finding that content-specific engagement opens doors generic use does not. Teachers who use AI for materials reach inclusion-specific use but those who use it only for administration do not. Breadth of engagement and professional distress are both at work along the gradient. In the capability-approach vocabulary (Sen 1999), three conversion factors – structural exposure (breadth), cognitive adjacency (materials) and motivational push (distress) – bridge the distance between AI adoption as a capability and inclusion-specific use as a functioning.

5.2. A cognitive boundary, supported by structural conditions

With GenAI, the domain boundary is not designed into the tool in the way it was with earlier systems. This is the theoretical point at which our findings meet the literature on post-adoption narrowing. The near-zero we observe in the coping item (3 of 69) cannot be attributed to technological limitation; the same tool is used daily by 51 of our respondents for other purposes. The category 'AI, as I use it' filters out the category 'AI, as it might help me here'.

Each of these claims rests on an identifiable combination of the two strands. The boundary itself is established quantitatively: the cascade (Figure 3), the conditional structure of the use-case fork (§4.3.1), and the near-absence of AI from the coping repertoire (§4.4). Its cognitive character is grounded qualitatively: Marina describes AI as generic where inclusion work is particular, and Camille can articulate an inclusion-specific next step she has not taken (§4.4). The exhausted trust of §4.2.2 draws on both: quality doubts concentrate among non-users with declining trajectories ($\rho? = -0.47$), and the open-ended material shows this distrust extending to the whole class of proposed solutions rather than to AI alone.

This is not a purely cognitive phenomenon. Structural conditions (§2.2) shape the cognitive environment in which teachers form a sense of what AI is for. Professional jurisdictions allocate inclusion to specialists (Abbott 1988), the division of labour within teaching keeps instruction and specialised support in separate compartments (Tardif and Lessard 1999), and the experience of activité empêchée (Clot 2008) reinforces relational rather than technical resources when inclusion difficulties arise. Some teachers do cross, through accumulated engagement under distress, or through anticipatory openness in its absence (§4.4.2), suggesting the boundary is bridgeable from more than one direction. Whether mainstream teachers consciously defer to specialist jurisdiction or simply do not conceive of AI as relevant to inclusion work remains an open question. Alvarez (2024) adds a further caution: assuming teachers should adapt to AI risks privileging those who can while marginalising those who cannot, reproducing within the profession the very exclusions that inclusive education seeks to address.

5.3. What does not predict AI use

Teachers with large classes, heavy inclusion burdens, prevented activity, and insufficient specialist support do not use AI more or less than teachers with lighter loads. A respondent with the highest unmet need in the sample articulates where her attention is directed:

'Specialist teachers do not have enough allocated periods to intervene effectively in classrooms (e.g. in my class there are only 45 minutes available per week!)' – Survey respondent, primary teacher, 15 years' experience, French-speaking Switzerland.

Her account is directed at specialist time, not technology. More AI tools or training are unlikely to bridge the cognitive category boundary by themselves. Teachers in inclusion distress already have AI in their hands; what is absent is a professional conception that makes AI legible as relevant to the problem. That absence is compounded by genuine limitations of current GenAI on relational, judgment-intensive tasks, but neither access nor training is the missing ingredient.

5.4. A global pattern

The pattern we document is not specifically Swiss. Across 55 economies, the OECD (2026) reports that support for students with special educational needs remains a minority AI application, and the same picture emerges in three other independent datasets on three continents (Lin, Cui, and Zhu 2026; Liu et al. 2026; Naatz and Ruppar 2025). We contribute by describing the within-teacher mechanism that sustains the gap among mainstream teachers, and by showing why generative AI changes the theoretical question rather than restating older ones.

5.5. Implications

We offer these implications tentatively. If the boundary is sustained cognitively and structurally rather than technologically, the intervention space is not 'more AI' but 'more bridging'. A tool designed with teachers (OECD 2026), beginning with content adaptation and gradually introducing student-specific functionality, may be more tractable than asking teachers to reconceive AI's role. Teacher education could treat the shift from content-specific to student-specific AI use as an explicit learning objective. The three teachers in the comprehensive subgroup (§4.3.2), their journeys, their accumulated use cases, their motivational profile, offer an existence proof that the boundary can be crossed. A policy agenda focused on bridging should study such cases and identify conditions under which comparable trajectories can be supported. Institutional scaffolding – peer exchanges among teachers who have crossed, case-based professional development – could make the passage visible to those who have not. The interpretive vocabulary used here (bridges, boundaries, pathways) is descriptive, not causal, only meant to organise an under-explored empirical space.

Any bridging agenda must also confront the ethical stakes of generative AI in inclusion work. Current systems can produce plausible but false advice, and a teacher acting on fabricated strategies for a vulnerable student faces consequences that a botched lesson plan does not. Inclusion-specific use raises privacy questions (describing a student's difficulties to a commercial system is categorically different from asking it to draft a worksheet), alongside risks of bias against the very learners inclusion is meant to serve, of diagnostic overreach by a tool with no clinical warrant, and of a dependence that erodes rather than builds professional judgment (Alhur et al. 2025). The strain of rapid, unsupported integration is itself documented as a source of technostress among educators (Khlaif et al. 2025). Furthermore, framing AI as a response to inclusion difficulties risks individualising what our data show to be structural failures: the teachers in greatest distress lack specialist time and institutional support, not technology (§5.3). AI may assist with adaptation, accessibility, and planning; it cannot substitute for relational judgment or adequate resources.

6. Conclusion

This article maps a specific gap in the encounter between generative AI and inclusive education. Swiss mainstream teachers adopt AI at rates that exceed international benchmarks, yet this adoption does not extend to the professional domain where they report the most distress. The cascade (74% use AI, 18% for inclusion, 4% as inclusion coping) is a steep disconnection at the coping layer. The variables that organise inclusion distress (need, workload, structural conditions) are largely disconnected from the variables that organise AI engagement (breadth of use, experience, school level). The two systems do not communicate.

GenAI has no built-in domain boundary. The only constraints on application are therefore the teacher's mental model of what the tool is for and the institutional conditions in which that model has been formed. The three conversion factors we identify – breadth of engagement, materials creation and distress – are candidates for intervention in professional development, tool design and institutional scaffolding, subject to confirmation in larger samples.

There are, however, several limitations to these findings: cross-sectional design, small convenience sample ($n = 69$), few observations at the upper levels of the gradient, and short-survey constructs that approximate rather than precisely measure the concepts. We hope that future work can explore the terrain we have mapped.

What Ertmer documented for computers in 2010, we tentatively document for AI in 2026. For many teachers, the first-order barrier is solved. The work now is to understand how teachers come to see (or fail to see) a tool they already have as relevant to a problem they already face, and how institutions can support that perception without reducing inclusive education itself to whatever the tool can do.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Notes on contributors

Roberta Besozzi holds a doctorate in social sciences and works as a research associate in the Department of Teacher Education at the University of Fribourg, Switzerland. Her research focuses on education and vocational education and training, professional pathways and identities, working and learning conditions, inclusive education, and artificial intelligence.

Anoop Menon holds a doctorate in Business Strategy from Harvard Business School and is co-founder of Eisengard AI. He previously served as assistant professor of strategy at the Wharton School, University of Pennsylvania. His research interests include strategy and strategic decision-making, cognitive neuroscience, decision theory and game theory, artificial intelligence, and the application of AI to complex social problems, including those arising from inclusive education.

ORCID

Roberta Besozzi http://orcid.org/0000-0003-4579-3289

References

Abbott, A. 1988. *The System of Professions: An Essay on the Division of Expert Labor*. Chicago: University of Chicago Press.

Alhur, A. A., Z. N. Khlaif, B. Hamamra, and E. Hussein. 2025. "Paradox of AI in Higher Education: Qualitative Inquiry into AI Dependency among Educators in Palestine." *JMIR Medical Education* 11:e74947–e74947. https://doi.org/10.2196/74947.

Alvarez, L. 2023. "Inclure grâce aux artefacts numériques: Des discours et des modèles pour penser." In *Éthique, numérique et idéologies*, edited by C. Roelens and C. Pélissier, 65–81. Paris: Presses des Mines.

Alvarez, L. 2024. *La numérisation de la formation et de l'éducation: Approche critique et interdisciplinaire*. Thèse d'habilitation, Université de Fribourg. https://doi.org/10.51363/unifr.lth.2024.035.

Alvarez, L., G. Ortoleva, D. Sutter Widmer, M. Fritz, J. Bugmann, S. Boechat-Heer, and C. Ramillon. 2024. "Future Teachers' Beliefs about Generative AI: Assessing Technology Acceptance as Students or as Aspiring Professionals." *Journal of Technology and Teacher Education* 32 (3): 383–408. https://doi.org/10.70725/379206cljimb.

Bandura, A. 1997. *Self-Efficacy: The Exercise of Control*. New York: Freeman.

Bešić, E., L. Schulz, and C. Schmid-Meier. 2025. "Inclusion and Equity in Artificial Intelligence (AI): Analysis of Educational Policies in Austria, Germany, and Switzerland." Preprint, *Research Square*. https://doi.org/10.21203/rs.3.rs-7232992/v1.

Besozzi, R., and A. Menon. n.d. *Odyssée professionnelle: Les enseignant·e·s face à l'inclusion scolaire*.

Capp, M. J. 2017. "The Effectiveness of Universal Design for Learning: A Meta-Analysis of Literature between 2013 and 2016." *International Journal of Inclusive Education* 21 (8): 791–807. https://doi.org/10.1080/13603116.2017.1325074.

Clot, Y. 2008. *Travail et pouvoir d'agir*. Paris: Presses Universitaires de France.

Creswell, J. W., and V. L. Plano Clark. 2018. *Designing and Conducting Mixed Methods Research*. 3rd ed. Thousand Oaks, CA: SAGE.

Cuban, L. 2001. *Oversold and Underused: Computers in the Classroom*. Cambridge, MA: Harvard University Press.

Davis, F. D. 1989. "Perceived Usefulness, Perceived Ease of use, and User Acceptance of Information Technology." *MIS Quarterly* 13 (3): 319–340. https://doi.org/10.2307/249008.

Ertmer, P. A. 1999. "Addressing First- and Second-Order Barriers to Change: Strategies for Technology Integration." *Educational Technology Research and Development* 47 (4): 47–61. https://doi.org/10.1007/bf02299597.

Ertmer, P. A., and A. T. Ottenbreit-Leftwich. 2010. "Teacher Technology Change: How Knowledge, Confidence, Beliefs, and Culture Intersect." *Journal of Research on Technology in Education* 42 (3): 255–284. https://doi.org/10.1080/15391523.2010.10782551.

Goodhue, D. L., and R. L. Thompson. 1995. "Task-technology Fit and Individual Performance." *MIS Quarterly* 19 (2): 213–236. https://doi.org/10.2307/249689.

Jasperson, J., P. E. Carter, and R. W. Zmud. 2005. "A Comprehensive Conceptualization of Post-Adoptive Behaviors Associated with Information Technology Enabled Work Systems." *MIS Quarterly* 29 (3): 525–557. https://doi.org/10.2307/25148694.

Katz, J. 2015. “Implementing the Three Block Model of Universal Design for Learning: Effects on Teachers’ Self-Efficacy, Stress, and Job Satisfaction in Inclusive Classrooms K-12.” *International Journal of Inclusive Education* 19 (1): 1–20. https://doi.org/10.1080/13603116.2014.881569.

Khlaif, Z. N., R. Alshakhshir, B. Hamamra, and A. Joma. 2025. “Reimagining Inclusive Education: The Assistive Power of Generative AI in Promoting Accessibility and Equity.” *British Journal of Visual Impairment*. https://doi.org/10.1177/02646196251382469.

Khlaif, Z. N., B. Hamamra, E. Bensalem, M. A. A. Mitwally, and M. Sanmugam. 2025. “Factors Influencing Educators’ Technostress While Using Generative AI: A Qualitative Study.” *Technology, Knowledge and Learning*. Advance online publication. https://doi.org/10.1007/s10758-025-09885-z.

Lin, X., H. Cui, and X. Zhu. 2026. “Chinese Special Education Teachers’ Experiences with Generative AI: Differences in Perceptions, Use, and TAM Pathways for IEP Writing and Lesson Planning.” *Journal of Special Education Technology*. Advance online publication. https://doi.org/10.1177/01626434261435251.

Liu, Ji, Airini Mbowe, Millicent Aziku, and Dahman Tahri. 2026. “Global Evidence on the Mediating Role of Teachers’ AI Beliefs in Linking AI-Related Professional Development and AI Infusion in Instruction.” *Education Sciences* 16 (4): 538. https://doi.org/10.3390/educsci16040538.

Miles, M. B., A. M. Huberman, and J. Saldaña. 2020. *Qualitative Data Analysis: A Methods Sourcebook*. 4th ed. Thousand Oaks, CA: SAGE.

Naatz, A. J., and A. L. Ruppar. 2025. “Special Education Teachers’ Use of Generative Artificial Intelligence (AI): An Exploratory Survey of Frequency and Factors Influencing Adoption.” *Journal of Special Education Technology* 41 (2): 227–239. https://doi.org/10.1177/01626434251379800.

Noy, C. 2008. “Sampling Knowledge: The Hermeneutics of Snowball Sampling in Qualitative Research.” *International Journal of Social Research Methodology* 11 (4): 327–344. https://doi.org/10.1080/13645570701401305.

OECD. 2026. *Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education*. Paris: OECD Publishing.

Okolo, C. M., and J. Diedrich. 2014. “Twenty-five Years Later: How Is Technology Used in the Education of Students with Disabilities? Results of a Statewide Study.” *Journal of Special Education Technology* 29 (1): 1–20. https://doi.org/10.1177/016264341402900101.

Sen, A. 1999. *Development as Freedom*. Oxford: Oxford University Press.

Tardif, M., and C. Lessard. 1999. *Le travail enseignant au quotidien: Expérience, interactions humaines et dilemmes professionnels*. Bruxelles: De Boeck.

Venkatesh, V., M. G. Morris, G. B. Davis, and F. D. Davis. 2003. “User Acceptance of Information Technology: Toward a Unified View.” *MIS Quarterly* 27 (3): 425–478. https://doi.org/10.2307/30036540.

Weber, H., A. Elsner, D. Wolf, M. Rohs, and M. Turner-Cmuchal. 2022. *Inclusive Digital Education*. Odense: European Agency for Special Needs and Inclusive Education.