



LUDIA: A Design and Evidence Statement

Designed to measure up

The goal: Universal Access to Universal Design for Learning.

LUDIA relaunches in August 2026. This document concerns the newly designed LUDIA, rebuilt from the ground up on a platform of its own. It is a different tool from the LUDIA that ran on the Poe platform from August 2023. Throughout, LUDIA means the relaunched tool unless the earlier one is named.

Beth Stark

Educator, Consultant, Inclusive Learning Systems Specialist

Jérémie Rostan

Middle School Principal, International School of Panama

Rights. © 2026 Beth Stark and Jérémie Rostan. This Design and Evidence Statement is licensed CC BY-NC-SA 4.0 (Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International). It may be shared, translated, quoted, and built on for non-commercial purposes, provided the authors are credited, changes are indicated, and anything built from it carries the same licence. LUDIA itself is released separately under its own licence, available at askludia.com.

Correspondence. hello@askludia.com

Suggested citation. Stark, B., & Rostan, J. (2026). *LUDIA: A design and evidence statement*. Ask LUDIA. <https://askludia.com>

Affiliations are given for identification only. This statement is the authors' own and does not represent the institutions named. August 2026.

This copy reflects a revision made in August 2026. We were unable to secure a satisfactory privacy confirmation from the model host we had initially intended to use, and LUDIA's model inference was moved to Google Vertex AI as a result. The privacy, data-governance, and transparency sections below describe that arrangement.

This paper draws on *Meet LUDIA, Your AI-Powered UDL Partner* (Rostan & Stark, 2023), published in *The International Educator*. The authors of this statement wrote that article, and it is quoted here with attribution. The article is freely readable and is not behind a paywall, but it is not published under an open licence.

Claims about privacy, accessibility, and sustainability summarise the statements published at askludia.com as they stood in August 2026. Those statements are maintained and will change. Where they differ from this document, they are the authoritative versions. Anyone relying on a claim of that kind should cite the version of this document they read and check the statement as it stands.

Contents

In brief	3
What LUDIA is designed to do	3
The Four T's: how LUDIA is meant to be used	6
The problem LUDIA was built to address	6
What changed, and why we rebuilt	9
Two versions, one goal	11
Where this begins	12
Why we are publishing this at relaunch	14
Evidence of Trust: a question we are keeping open	15
Part One: The EdTech for Good pillars	17
Part Two: The European Commission's guiding questions	21
Part Three: The EdTech Quality Indicators	27
Part Four: What we will count as evidence	30
One thing, well	33
Designing for sustainability	34
What happens next	35
A closing note on trust	36
Funding and competing interests	36
Citation and references	37

In brief

LUDIA is a no-cost, private, multilingual AI thought partner that connects educators with the Universal Design for Learning framework. It relaunches in August 2026 on a platform of its own, purpose-built for LUDIA, replacing the version that ran on Poe from 2023.

This paper reads the relaunched LUDIA against four guides published in 2026: the United Nations Children's Fund (UNICEF) EdTech for Good Framework, Version 2.0, the European Commission's guidelines on the ethical use of AI in teaching and learning, the EdTech Quality Indicators Guide, and the European EdTech Alliance's Needs-Based EdTech Evidence Mapping. It sets out what has been built and can be checked, what cannot yet be shown, what we intend to gather from August 2026, and what LUDIA will not become.

We hold no evidence that LUDIA improves learning, and we say so throughout. We say the same about the evidence base of the framework LUDIA rests on, which is under active and legitimate challenge. This is a design and evidence statement, not a marketing document and not a claim of impact.

What LUDIA is designed to do

No cost, safe, private, and universally accessible. A multilingual AI thought partner built to connect with the Universal Design for Learning (UDL) framework. Any general-purpose model can be prompted towards inclusive design if an educator, instructional designer, or learning leader knows what to ask for. LUDIA is engineered to do it by default, in every exchange, without being asked, redirecting educators away from fixing the learner and towards reducing barriers in the design of instruction and learning environments. The barrier lives in the design, not in the learner. Universal Access to Universal Design for Learning. That is the goal.

LUDIA is built for educators, instructional designers, and learning leaders. Where this document says educators on its own, it means all three unless the context narrows it.

What the UDL Guidelines are

Within schools and learning organisations, defaulting to a fixed one-size-fits-all approach shortchanges learners, especially those furthest from the mythical average learner it was designed for. CAST, the nonprofit education research and development organisation that created the UDL Guidelines, has been working on this since 1984, and Universal Design for Learning has been a named framework since the 1990s. Developed by David H. Rose, Anne Meyer, and colleagues at CAST, the framework supports educators, instructional designers, and learning leaders as they design learning experiences and environments for the full range of learners in their care, anticipating and embracing variability from the start, rather than focusing on perceived deficits in learners. CAST states the premise plainly: "there is no average brain. Variability is the dominant feature of the nervous system" (CAST, 2018). CAST

describes the three principles, Engagement, Representation, and Action and Expression, as aligning to three broad networks in the learning brain: affective, recognition, and strategic. CAST is careful about this, calling it a simplified model that is useful for thinking about design rather than a claim about how the brain works (CAST, n.d.). We return below to how well that framing is supported. The CAST UDL Guidelines translate those principles into concrete design considerations to reduce barriers to learning.

The UDL Guidelines are not a prescription or a checklist. CAST describes them as a tool offering suggestions that can be applied to instructional design to reduce barriers, honour learners' multiple identities, and maximise learning opportunities according to specific learning goals (CAST, 2024). They do not say which method to teach by. What they give educators and instructional designers is a lens for looking at what is already in place and finding the barriers in it, which is why UDL sits alongside a wide range of approaches.

How LUDIA works with the Guidelines

None of which makes UDL neutral, and we would not claim that it is. UDL rests on a position about learners, that variability is the rule and not the exception, and the UDL Guidelines 3.0 went further still, taking on barriers rooted in bias and in systems of exclusion. CAST also states that Guidelines 3.0 seeks to put UDL in conversation with other asset-based approaches and theoretical frameworks, naming Funds of Knowledge and Disability Studies among them, as a partner in elevating learner variability within those pedagogies (CAST, 2024). That is the pairing LUDIA is built on: the Guidelines themselves, and the asset-based pedagogies and practices they were designed to work alongside. Which method to teach by is the educator's or instructional designer's decision.

LUDIA is there at the moment that matters most, while a design decision is still open, and it is just as useful afterwards. It supports educators in anticipating who a design might exclude, and why, working through engagement, representation, and action and expression before a learner meets the barrier. It is equally suited to reflection, when a lesson or a learning environment is not landing the way an educator hoped: brought the question "why isn't this working?", LUDIA helps trace the answer back to a barrier in the design, rather than to the learner. Throughout, it holds learner agency, and the conditions that foster it, at the centre. And it connects the particular situation at hand to the relevant UDL Guidelines, so that the reasoning is visible and not only the suggestion.

Most AI agents are engineered to be solution engines: deliver an answer, give the user the output they came for. LUDIA is engineered to be a thought partner. LUDIA can draft the kinds of planning materials any other AI agent can draft, but that is not the goal. LUDIA is purpose-built to help educators, instructional designers, and learning leaders understand and make connections to the UDL Guidelines, one situation at a time. By grounding every response in the framework and its Guidelines, LUDIA redirects attention to barriers in the design, not in the learners, and supports educators in expanding their UDL practice in context.

Designing against the defaults

The name LUDIA is a remix of the acronyms AI and UDL, and a play on ludum, the Latin word for play and learning. The first LUDIA took the pronouns they and them, chosen to draw attention to a documented default. In an analysis of 1,375 chatbots listed on chatbots.org, 63.6 percent carried at least one gender cue in their name, avatar, or description, and where a cue was present it was overwhelmingly female: 76.9 percent of the gendered names, 77.6 percent of the gendered avatars, and 67.4 percent of the gendered descriptions. The authors call this a gender bias in chatbot design and recommend against female-by-default designs (Feine et al., 2020), a finding consistent with UNESCO's account of voice assistants designed female or female by default (West et al., 2019).

The relaunched LUDIA takes the pronoun it. The EU Artificial Intelligence Act (AI Act), Regulation (EU) 2024/1689, requires at Article 50 that users be told they are interacting with an AI system, and the European Commission's guidelines put the responsibility for that disclosure on the people providing a tool rather than on the educator using it. A human pronoun works against both, and so does a human face. LUDIA is drawn as a cartoon chatbot with headphones and a speech bubble, and the drawing stops well short of a human likeness.

There is a related tendency worth naming, because it is the one most likely to defeat what LUDIA is for. Models trained on human feedback often drift towards sycophancy: they match or validate the user's stated view instead of challenging it. Sharma et al. (2023) found this across five leading AI assistants, in tasks where the models wrongly conceded mistakes, gave biased feedback, and repeated users' errors back to them. They traced it to the training itself, since matching a user's view turns out to be one of the strongest predictors of which response a human rater will prefer. A tool designed to validate a user's perspective is pleasant company. A tool designed to identify and reduce barriers has to be willing to say something the user did not want to hear. That is a statement of intent, and nothing more. The tendency Sharma describes belongs to the training of the general-purpose model LUDIA runs on, and no instruction we write removes it. We hold nothing showing that LUDIA resists it. Part Four sets out how we mean to find out.

We intentionally declined design choices that are typically aimed at bringing users back. There are no streaks, no badges, and no attempt to make the exchange feel like a friendship with a fellow human. The in-app notices that do exist are functional, confirming an upload or a download, and an educator controls where they appear and how long they stay. What LUDIA offers instead is a slow, continuing professional learning exchange about Universal Design for Learning, specific to the situation an educator brings to it. LUDIA is grounded in three values: designing for access, designing for agency, and amplifying the expertise of educators, instructional designers, and learning leaders.

The Four T's: how LUDIA is meant to be used

In 2023 we first published a way of engaging with LUDIA that we call the Four T's. The guide is published at no cost, and it predates every guide discussed below.

TELL Tell LUDIA about the learners in their care, the learning plans, and the challenges. The more thorough the telling, the better.

TINKER Continue the chat. Ask for clarification, examples, developments, and materials, and explore the options and strategies LUDIA shares.

TWEAK Consider the provisional outcomes of the exchange. Adapt them as fits the context, and as the work proceeds.

TRANSFER Reflect on the impact that identifying and reducing barriers brings to learners and to the learning environment. Connect, consolidate, and deepen that learning through Visible Thinking Routines such as "I Used to Think, but Now I Think", "Connect, Extend, Challenge", or "3-2-1 Bridge".

The first response is not where the value lies. Transfer is the step where an educator reflects on what identifying and reducing a barrier brought to learners and to the learning environment, and then connects, consolidates, and deepens that learning. The routines named there are drawn from Project Zero's Thinking Routines Toolbox, a collection of many such routines. An exchange that stops at the second T is LUDIA serving a purpose, but LUDIA is not yet being used fully as a thought partner.

The problem LUDIA was built to address

We first wrote about LUDIA in *The International Educator* in September 2023, and what we said then still holds. The "wow factor" of AI was undeniable, we wrote, and many educators had already tapped into the endless ways it can increase efficiency. But "AI integration will not be truly effective and impactful in the world of education unless it helps us to augment and amplify excellent pedagogical practices" (Rostan & Stark, 2023). We started with a question, not a technology: could AI help create the conditions for real agency in learning? UDL is both a method and a lens. It rests on the science of how people learn, and it helps educators and learning leaders find and remove barriers without lowering what is expected of anyone (Meyer et al., 2014).

UDL's standing in law and policy

UDL's standing in international human rights law comes through interpretation rather than treaty text. The United Nations Convention on the Rights of Persons with Disabilities (CRPD) establishes the right to inclusive education in Article 24, and Article 2 defines universal design as it applies to products, environments, programmes, and services. Neither article

uses the phrase universal design for learning; that phrase appears only in the CRPD Committee's General Comment No. 4 (2016), the authoritative interpretation of Article 24, which directs States to adopt the universal design for learning approach and describes it as providing teachers and other staff with a structure for creating adaptable learning environments that meet the diverse needs of all learners. The commitment to inclusive education that Article 24 rests on goes back further still, to the UNESCO Salamanca Statement (UNESCO, 1994) and Sustainable Development Goal 4, though neither of those, either, names UDL directly.

UDL's clearer standing is in national law: four instruments make the point without needing a survey behind them. In the United States, the Higher Education Opportunity Act of 2008 defines universal design for learning in statute, and the Every Student Succeeds Act of 2015 adopts that definition. In Chile, Decreto 83 of 2015 sets the criteria for curricular adaptation in early years and basic education and builds them on the *Diseño Universal para el Aprendizaje*. In Colombia, Decreto 1421 of 2017 requires schools to incorporate inclusive education and the *diseño universal de los aprendizajes* into their institutional education plan. In Spain, the LOMLOE, Ley Orgánica 3/2020, places the application of DUA principles among the principles and aims of education, and the curriculum decrees made under it carry it into each stage. Beyond these four, UDL appears in curriculum guidance, higher education policy, and disability frameworks in a number of other countries. We have read some of those documents. We have not compiled them to a standard we would ask anyone to rely on, so we give no count here. UNESCO's 2020 Global Education Monitoring Report puts the expectation directly: all teachers should be prepared to teach all students, and inclusive approaches belong at the core of teacher education rather than treated as a specialist topic (UNESCO, 2020).

The scale of the gap

Law and policy affirm UDL's standing; what has not followed at the same pace is the professional learning educators need to enact it in their own context. The global EdTech market is projected to reach four hundred and thirty-seven billion US dollars by 2033, on a forecast published by Precedence Research, a market-research firm. We use it only to indicate the scale of spending. That spending has not closed the inclusion gap. Globally, 240 million children are experiencing disability, and roughly half of those in low- and middle-income countries are out of school (UNICEF, n.d.). Children with visual, hearing, physical, or intellectual disabilities are 2.5 times more likely not to attend school at all (UNESCO, 2020). That gap has structural drivers: inaccessible infrastructure, discriminatory attitudes, and under-resourced schools, not a limitation in the learner. Learners with exceptional strengths are also underserved: international research finds that up to 50 percent of high-ability learners underachieve at some point during schooling (Siegle, 2018), and, in one study of upper-percentile secondary students in Hong Kong, 53 percent (Phillipson, 2008). For these learners, the driver is rarely inability: research points instead to boredom, social-emotional isolation, and a curriculum that may not feel meaningful or challenging. Learners who hold both exceptional strengths and disabilities are frequently missed by identification systems built to look for one or the other (Foley Nicpon et al., 2011). Across OECD systems, 33 percent

of teachers identify a shortage of competencies for supporting students with specific needs (OECD, 2024).

Those figures come from different populations, methods, and years, and they do not add up to one measurement. The Hong Kong figure is a finding about one sample of upper-percentile secondary students, not a global rate. We set them side by side because each shows part of the same gap. No one of them establishes it. The learners named here, disabled, gifted, twice-exceptional, are illustrations of variability, not the limit of it. UDL does not treat them as separate populations requiring separate accommodations; its premise is that variability is the rule for every learner, and that a design built for one mythical average learner will always leave someone out. The scale of global EdTech investment is beyond debate. The question worth asking is whether that investment has improved educators' ability to design learning experiences and environments through the lens of inclusion. None of this is a case for LUDIA on the strength of the numbers alone. But the numbers point to where leverage actually sits: not in another platform or another billion dollars of spending, but in the moment an educator or instructional designer makes a design decision. That is a small point of intervention against a structural problem, and it is also the one point where a barrier can be identified and reduced before a learner ever meets it. LUDIA was designed for that moment.

Barriers educators face in practice

Educators and instructional designers face hurdles in developing their capacity as UDL implementers. The responses we hear most often, compiled and shared here from years of workshops and coaching conversations rather than quoted from any one person, include: "UDL feels like one more thing, and I am not even sure where to start. I already have a lot of demands on my time. It sounds great, but how do you do Universal Design for Learning?" Responses like these are common and, we have to say, valid. We named the obstacle in 2023 as "a trifecta of barriers" that "can deter us from deepening our understanding and utilization of the principles of UDL: context, complexity, and our readiness to leap over the great 'knowing-doing' divide" (Rostan & Stark, 2023).

UDL guidance exists in abundance, but it is necessarily written in general terms, while a barrier is always particular to this learning engagement, these learners, this room, this week. LUDIA was built to bridge this divide.

What we committed to

Two commitments came out of that, and we made both in 2023, before any of the guides discussed below existed. The first concerns mindset. Partnering with LUDIA gives educators a new way to think through, adapt, and plan for challenges that create barriers for learners in their care, and it reinforces the key mindset that barriers do not originate in the learner. When we focus on strengths and embrace learner variability, we create space to identify and reduce barriers in the design of instruction and the learning environment. LUDIA can support educators in replacing feelings of hopelessness and helplessness with a sense of what is possible for all learners, and the adaptable barrier-reducing options LUDIA offers are in response to their unique context and within their circle of influence as designers.

The second commitment concerns AI, and we would put it today exactly as we put it then:

AI can certainly enhance but not replace our thinking as educators.

Rostan and Stark (2023), Meet LUDIA, Your AI-Powered UDL Partner, The International Educator, 27 September 2023

That article also set out what we did not want LUDIA to become, and the wording we used then is the wording we would use now. LUDIA "invites educators to envision their problems of practice, not so much as problems, but as puzzles to solve." Put plainly: LUDIA is not designed to be a problem solver. The connections, we wrote, "happen in the creative back-and-forth, the subsequent tinkering and tweaking, evaluating, and adapting for learners and the unique context in question," through which "educators can grow through professional play" and develop their competencies as learning designers. We also stated plainly that "LUDIA retains all the limitations of generative AI" and that "the richest learning happens when LUDIA's responses serve as scaffolds and entry points" (Rostan & Stark, 2023). None of this was written to satisfy a framework. That is why the guides discussed below fit as well as they do, and where one of them asks for something we did not consider in 2023, we say so.

What changed, and why we rebuilt

LUDIA was first introduced in August 2023 on Poe, Quora's Platform for Open Exploration. Poe had made the power of complex technology accessible for creators and users alike, and it allowed two educators to put a working UDL thought partner in front of colleagues without building infrastructure. It worked. According to traffic data from the Smore page built for LUDIA, the tool has reached visitors in more than 120 countries. We report this figure as the platform provided it and cannot independently confirm it.

In August 2026, a newly designed LUDIA relaunches on a platform of its own, built from the ground up for safety, privacy, and accessibility. The goal has not changed. What changed is that a set of commitments we had been making could not be guaranteed on a third-party commercial platform, and in the end that weighed more heavily than the convenience of not building.

Part of what changed was something we observed. We offer what follows as our experience of running LUDIA on a commercial platform between 2023 and 2026, not as a documented finding, and we did not keep measurements. Over those three years, two things seemed to us to move in the same direction. LUDIA's responses drifted. The ceiling on what an educator could do without paying came down. Neither shift was ours to control. We are not attributing bad faith to anyone: a platform opens generously, attracts the people who will build on it, and then, under its own commercial pressures, what is available at no cost narrows. For most products this is an irritation for consumers. For an educator working where resources are scarce, they may experience these shifts as a door closing.

ACCESS On Poe, reaching LUDIA meant creating an account on a commercial platform first, and then working within whatever that platform currently allowed a non-paying user to do. For a tool whose entire purpose is reducing barriers, an account requirement is a barrier, and a moving allowance is a barrier that moves. Both fall hardest on the educators with the least institutional backing. The relaunched LUDIA requires no account, no email address, no institutional login, and there is no allowance to run down.

ACCESSIBILITY We could not promise a conformance standard for an interface we did not control. Having a front end of our own is what makes it possible to design to the Web Content Accessibility Guidelines (WCAG) 2.2 Level AA and to publish an accessibility statement describing what that means in practice: keyboard-only operation, screen reader support, browser zoom to 400 percent with content reflowing, light and dark colour schemes, reduced-motion preferences honoured, a visible focus outline on every control, and tables that become stacked cards on small screens so that nothing is cut off. For a tool built on UDL, inheriting someone else's interface was the wrong compromise.

PRIVACY Accessibility and privacy are not settings in LUDIA. They are the architecture. The relaunched build sets no cookies, requires no account, writes no chat content to our servers, and collects no data for model training. Educators own their exchanges, and any conversation can be downloaded and uploaded again later to continue where it left off. On a third-party platform, what was retained and the reasons why were not ours to determine.

LANGUAGE LUDIA on Poe could respond in a language other than English, based on the browser's detected language or on the language a question was asked in, but this was never offered as a dedicated feature, and consistency across a conversation was not guaranteed. The relaunched LUDIA offers 13 languages as dedicated options, with more to follow, so an educator can choose to work in one of them rather than rely on the tool inferring it. LUDIA is authored in English and the other languages are machine-translated. They have not been validated by professional translators, and they may contain errors. That is a limitation of our resources, not of those languages, and we share this openly.

INDEPENDENCE LUDIA is an open-ethos resource designed to support educators in making progress towards Sustainable Development Goal 4 of the United Nations: inclusive and equitable educational opportunities for all. There is no premium tier, no advertising, no sponsored content, no commercial referrals, no data collection for model training, and no commercial relationship with curriculum vendors. That leaves an obvious question, and it deserves a response. Hosting and model inference are not free, and costs are currently manageable. Costs are manageable because the architecture was designed to keep them so, not by accident. The first step in our development plan was to establish that a tool of this kind can exist on these terms at all, and the only way to establish it was to have it built and to pay for it. The architecture keeps costs manageable, and growing LUDIA's reach will not change this aspect of the design. When LUDIA scales, access stays at no cost to the educator, private by architecture, and accessible by design. Cost-free is not the same as always-available, and the goal on the cover is not a guarantee of service. The same self-funded architecture that keeps LUDIA free limits how many educators it can serve at one moment and caps what can be spent running it in a month. Both limits are set out under Technical robustness and safety,

and both are real. We can hold the price and the terms of access. We cannot yet promise the capacity behind them. LUDIA is not a revenue-generating product, we are not a start-up, and we are not hoping LUDIA will someday be bought out. LUDIA was created for educators, by educators. Owning and overseeing the platform is what makes those commitments ours to keep.

The relaunched LUDIA is new: new code, new interface, new languages, new behaviour. Whatever confidence the earlier version earned over three years attaches to a tool that no longer exists in that form. We are beginning the evidence again, and the remainder of this document is written on that basis.

Two versions, one goal

The two are easily confused, and nearly everything written about LUDIA so far describes the first. Here is the difference set out directly.

Table 1. LUDIA on the Poe platform, August 2023 to August 2026, compared with the relaunched LUDIA from August 2026.

Aspect	LUDIA on Poe, August 2023 to August 2026	LUDIA from August 2026
Where it runs	Poe, a third-party commercial platform	A platform of its own, built from the ground up for safety, privacy, and accessibility
Getting in	Required a Poe account	No account, no email address, no institutional login, nothing to install
Accessibility	Inherited the host platform's interface	Designed to WCAG 2.2 Level AA, with a published accessibility statement
Privacy	Retention determined by the host platform	Retention is determined by how the host institution configures its Google Cloud project. No cookies are set, no chat content is stored server-side, and no data is collected for model training; exchanges are downloadable and re-uploadable. Google does not use prompts or responses to train or fine-tune its models, and this deployment is not subject to prompt logging for abuse monitoring; see Privacy and data governance

Aspect	LUDIA on Poe, August 2023 to August 2026	LUDIA from August 2026
Languages	Responded in the browser's detected language, or in whichever language the question was asked in; not offered as a dedicated set of languages	13 languages and growing, offered as dedicated options; translations other than English are machine-generated
Cost and model	No cost to the educator, within limits set and revised by the host platform	No cost to users. No premium tier, no advertising, no sponsored content, no commercial relationship with curriculum vendors
Evidence base	Referenced independently in research, articles, and podcasts; no study of its effects	None yet. Beta tested by volunteer educators internationally ahead of relaunch; documenting and sharing what that testing found is part of our plan
Goal	Universal Access to Universal Design for Learning	Unchanged

Throughout the remainder of this document, LUDIA means the version relaunching in August 2026. Where something applies only to the earlier tool, it is named as such.

Where this begins

LUDIA is designed to public-good terms, not a consumer product, in the spirit of UNESCO's Windhoek+30 Declaration, which names information as a public good (UNESCO, 2021b). Next to it sit UNESCO's Recommendation on the Ethics of Artificial Intelligence (UNESCO, 2021a), which lists the values any use of AI in education should be weighed against, including human dignity, inclusion, and a flourishing environment, and its Guidance for Generative AI in Education and Research (UNESCO, 2023), which asks for an approach that keeps people at the centre. As noted above, UNESCO's 2020 Global Education Monitoring Report says the underlying expectation plainly: all teachers should be prepared to teach all students, and inclusion belongs at the heart of teacher education rather than off in a specialist corner. Those are the commitments that shape how we design LUDIA.

In 2026, four further guides were published, each asking in its own way whether a digital tool belongs in education. We are using all four to inform our process, not to grade what we have designed. They are a set of questions we intend to keep answering, because designers are

never finished, and because progress is both nuanced and non-linear. Safety, contextual fit, accessibility, who gets left out, and whether a tool can be trusted at all are forms of evidence too, and these four guides regard them as such.

EdTech for Good Framework, Version 2.0. United Nations Children's Fund (UNICEF), Helsinki, 2026. Licensed CC BY 4.0. Asks who is behind a product, what it is intended to do, and whether it is appropriate for responsible use, the last through five pillars: Safety and Well-being, Educational Impact, Designed for Teaching and Learning, Contextual Readiness, and Inclusive Access. Informed by feedback from more than 140 organisations across over 60 countries.

Guidelines on the ethical use of artificial intelligence and data in teaching and learning for educators, revised edition. European Commission, Directorate-General for Education, Youth, Sport and Culture. Publications Office of the European Union, Luxembourg, 2026. <https://doi.org/10.2766/7967834>. Licensed CC BY 4.0. Puts guiding questions directly into the hands of teachers and school leaders across eight requirement areas, and situates them within the AI Act and the General Data Protection Regulation (GDPR). Non-binding, and explicit that it does not provide enforcement guidance.

EdTech Quality Indicators Guide. EdTech Quality Collaborative, June 2026. Licensed CC BY-NC-ND 4.0, with attribution to the EdTech Quality Collaborative. Evaluates education technology against five indicators: Safe, Evidence-Based, Inclusive, Usable, and Interoperable. Developed by 1EdTech, CAST, CoSN, Digital Promise, ISTE+ASCD, and SETDA.

Needs-Based EdTech Evidence Mapping. Lindroos Cermakova, A., Havinga, B., & Clary, A. European EdTech Alliance, Bielefeld, March 2026. Licensed CC BY-NC-SA 4.0. Sorts evidence into five domains according to the purpose it serves rather than its methodological strength, and argues that the sector should move from proof of impact towards proof of trust.

Each one gives us something the others do not. UNICEF speaks from child rights, and asks who is behind a product and whether it is fit to be used at all. The European Commission puts the questions in an educator's hands and sets them next to binding law. The EdTech Quality Indicators Guide matters to us for a particular reason: CAST is where UDL came from, and Digital Promise has shaped much of the practice LUDIA points people towards, so a tool built on UDL should be willing to be read against a guide CAST helped develop. The Needs-Based Evidence Mapping is not an evaluation framework at all. We regard it as a way of asking what evidence is for, and it organises the last part of this document.

Why we are publishing this at relaunch

Usually a guide like this comes out after the fact, once a tool is established and someone needs it justified. Used that way, a guide becomes a scoring exercise, because the honest answer to most of its questions was settled by decisions taken months earlier. We would rather use them the other way round: as design guidance while the design can still change, and as a statement, made in advance, of what we intend to count as evidence.

The line under the title should be read as a verb and not a verdict. We are not claiming that LUDIA measures up. We are saying that it is designed to, and that designing to is continuous work. This document is where we set down what we have designed, what we have not, and what we intend to count as evidence. It is fixed at a moment, dated August 2026, so a future revision can be measured against it. That revision, whenever it comes, is not a verdict either, and the distance between the two is the point.

A relaunch is a good moment to start. We made these commitments while the new LUDIA was being developed, not afterwards to fit what came out, and the gaps named here are ones we are naming ourselves. This document describes what we have designed. Almost all of it can be verified by using LUDIA. It says plainly what we cannot yet answer. And it says in advance what we intend to gather, and against which guide, so that educators and learning leaders can hold us to it later.

A note on context: the four frameworks were written for different purposes and from different vantage points. The EdTech for Good Framework is a global, child-rights instrument. The European Commission's guidelines are non-binding guidance for teachers and school leaders, sitting alongside binding law. The EdTech Quality Indicators Guide was developed by US-based organisations, mainly with K-12 schools in mind. The Needs-Based Evidence Mapping is a European ecosystem study concerned with how evidence travels between educators, providers, policymakers, and investors. Any judgement against any of them depends on context and shifts as the field changes. It also follows that not every indicator in every guide is aimed at a tool like this one. Where an indicator assumes a category LUDIA does not belong to, we say so and say what a reader gives up by our not meeting it, rather than record it as a failure or quietly reframe it as a strength. What follows is our own reading and our own commitment. None of the authors of these publications has reviewed or endorsed LUDIA. Nothing here is legal advice, and our reading of the AI Act is our reading, not a determination. We will revisit this document as the relaunched tool is used and as the field develops.

Two limits belong to this document as a whole, and both are easier to see gathered than scattered. The first is that nothing here has been independently verified. Nobody outside has audited LUDIA's accessibility, its security, or its reliability. No independent researcher has reviewed the evidence claims made in it. The usability testing was informal and undocumented. The Accessibility Conformance Report is in preparation and is self-authored, which is normal for such reports and is not an audit. Even this document's conformance as a tagged PDF has been built to the standard and not confirmed by an external checker. Where

we engage someone to build or check part of LUDIA, that is professional practice and we are glad of it. It is not independent verification, and we will not present it as such. Each appears in its own section below, at the point where it belongs. Together they say something the separate disclosures do not: a reader has our word and the tool itself, and nothing else. The tool is the part that can be checked without us, which is why so much of this document points at what anyone can verify by using it.

The second is that we chose these four guides. Nobody assigned them, no procurement process required them, and we selected them after LUDIA was largely built. Our reasons are given for each, and one of the four was co-developed by an organisation that gave us an award, which is disclosed under competing interests. There are other instruments we did not use, including product certifications and privacy evaluations that require submitting a tool for external review rather than reading it against a published rubric. We have not sought those. Reading yourself against guides you picked is the weakest form of assessment there is, and it is the form two people can carry out with the resources they have. We would be glad to see someone read LUDIA against a guide we did not choose.

Evidence of Trust: a question we are keeping open

The Needs-Based Evidence Mapping proposes that EdTech shifts its centre of gravity away from an exclusive focus on proof of impact and towards a broader culture of proof of trust: evidence that is credible, context-rich, explainable, and socially responsible. We find the framing useful, and we have taken it as the organising idea of this document and our future inquiry.

Our reason for keeping it open comes from the report itself. Writing about learning impact, the report's authors point out that the effects educators value most are often the ones evaluation is least equipped to capture: whether a classroom works together differently, whether the exchanges between an educator and a learner have changed in quality, whether a particular learner can take part at all. Effects of this kind, the report notes, remain "methodologically underexamined", because they resist both experimental design and simple measurement. The conclusion the authors draw is the one we want to carry forward:

"Learning impact" cannot be treated as a singular, universally agreed construct.

Lindroos Cermakova et al. (2026), Needs-Based EdTech Evidence Mapping, Section 2.1

If learning impact cannot be treated as a single agreed construct, neither can trust. Trust is not a property a tool possesses and can be certified as holding. Trust sits between a person, a tool, and a setting, and has to be earned again whenever any of the three changes. A guide can make the grounds for trust visible. It cannot manufacture the trust itself. The report's own

caution about how evidence mutates as it travels applies here with some force: "proof of trust" is just the kind of phrase that ends up as a badge, both an actual certification mark and a metaphorical claim to have settled the question.

The four guides bear this out, because each defines trust differently. For UNICEF it is a function of transparency: whether enough has been disclosed about a product and the people who created it for a judgement to be possible at all. For the Needs-Based Evidence Mapping it is the coherence and honesty of an evidence portfolio. For the Quality Indicators it is largely a matter of safety and privacy protections. The European Commission offers a fourth account, describing a trustworthy AI tool as one that is consistently reliable, transparent about how it works, respects privacy, avoids bias, and supports learning in ways that align with the values of the school community. That last clause is the one we keep returning to, and the guidelines make its implication explicit:

the context of use is key for ethical decision making

European Commission (2026), Guidelines on the ethical use of artificial intelligence and data in teaching and learning for educators, p. 21

If a learning organisation's values are part of what makes a tool trustworthy, no self-assessment can settle it. We do not know the school or institution. That matters for LUDIA in a particular way. LUDIA works with the educator or instructional designer, upstream of the learner. The most it can change is a design decision, and what a design decision changes is a barrier. Those are the relational, contextual effects the evidence report calls valued and underexamined, and no launch metric will catch them. This is not modesty about our evidence. It is a statement of what we will be looking for as we go on.

Part One: The EdTech for Good pillars

The guide. EdTech for Good Framework, Version 2.0. United Nations Children's Fund (UNICEF), Helsinki, 2026. Licensed CC BY 4.0. Five pillars: Safety and Well-being, Educational Impact, Designed for Teaching and Learning, Contextual Readiness, and Inclusive Access.

The EdTech for Good Framework asks its questions in three layers, and it is explicit that early-stage tools are expected to answer them differently from mature ones. Basic evidence, in its terms, shows that an organisation can explain its approach and provide an acceptable degree of documented transparency. That is where the relaunched LUDIA sits, and it is the tier we are aiming at here. Two of the three layers are about context, not judgement, and we answer those briefly.

Two educators co-created LUDIA, and designed it to public-good terms rather than a company's product. There are no investors, no advertisers, and nobody making money from data. The only thing we depend on from outside is the model LUDIA runs on, which we name and can change. On the framework's question about expertise, we bring decades of collective international teaching and leadership experience, teaching students who are blind and have low vision, IAAP certification as a Certified Professional in Accessibility Core Competencies (CPACC), and years of coaching and programme design with international schools. As creators of LUDIA, we were honoured to receive CAST's Anne Meyer UDL Design Award in 2025.

Safety and well-being first

DESIGN COMMITMENT LUDIA is kept safe by the way it is built, not by a rulebook around it. LUDIA is for adult professionals. It is not directed at children, and entry requires the user to confirm they are 18 or older. That confirmation is a self-declaration and not a verification. It establishes who LUDIA is for, not who reaches it. The protection is not the checkbox. There are no accounts, no profiles, and no chat content written to our servers, so LUDIA stores nothing about any user, of any age. LUDIA says it is an AI chatbot at the point of use. It suggests. It does not decide, and nothing it produces reaches a learner except through an educator who has judged it worth using. There are no streaks and no badges, and nothing in the visual design or the tone invites an educator to mistake LUDIA for a person. When the tab closes, the session is over. A Data Processing Agreement is in place with the model provider under Google's standard Cloud Data Processing Addendum, covering the processing carried out on Vertex AI.

NOT YET IN PLACE A public process for what happens when we change the model LUDIA runs on is being drafted and is not yet published. It will live in the versioned behavioural specification described under transparency and explainability. What is in place is a documented incident-response plan, with 72-hour breach notification to the supervisory authority under Article 33 of the GDPR, and a completed Data Protection Impact Assessment.

LUDIA is currently operated by Beth Stark as data controller while a formal legal entity is registered in Germany, and the privacy statement will be updated on incorporation.

Educational impact

DESIGN COMMITMENT LUDIA is not an authority. It does not grade, predict, rank, or automate any educational decision. Its responses are options to weigh, not instructions to follow. The unintended impact we are designing against is deskilling, since a tool that answers UDL questions can quietly become a substitute for thinking about UDL. LUDIA mitigates this by pointing back to CAST's framework and to sources, instead of treating its own output as the destination.

NOT YET IN PLACE Everything else. The relaunched LUDIA has no outcome data, no efficacy trial, and no monitoring, because it has not yet been used. Our intended claim is narrow and remains untested: that a UDL-literate thought partner, available at the moment a lesson or a policy is being designed, helps an educator or instructional designer notice a barrier they had not seen and consider a design option they had not considered.

Designed for teaching and learning

DESIGN COMMITMENT LUDIA carries no teaching method of its own. CAST is explicit that the Guidelines are not a prescription or a checklist, and they do not say which method to teach by. CAST is equally explicit that Guidelines 3.0 seeks to put UDL in conversation with other asset-based approaches, naming Funds of Knowledge and Disability Studies among them, as a partner within those pedagogies (CAST, 2024). So what LUDIA is grounded in is the UDL Guidelines, published and maintained by CAST, together with the asset-based pedagogies and practices they were designed to work alongside. LUDIA connects educators to both and leaves the teaching decisions where they belong. The inline links to relevant resources and next steps in every response are a deliberate professional-learning choice, to create a UDL-informed experience that invites educators or instructional designers to explore more.

IN PROGRESS This pillar also asks about participation and co-creation. Ahead of relaunch, a global community of volunteer educators beta tested LUDIA in their own planning and design work, across languages, education systems, and cultural contexts. That is the most substantial participation in LUDIA's history, and it shaped the rebuild directly. Documenting and sharing it is part of our plan, and it is under way and not finished: who took part, what they told us, what changed as a result, and what we chose not to change. When we first wrote about LUDIA we also said that beyond this tool, we envisioned a movement by creating online spaces for educators to share their stories, successes, and recommendations. We look forward to designing this in the future.

Contextual readiness

DESIGN COMMITMENT LUDIA requires a browser and a connection, and nothing else: no procurement, no installation, no administrator, no per-seat licence, and no integration. It relaunched in 13 languages. Two capabilities in development speak directly to this pillar: a

plain language version prioritising the UDL Access Guidelines, and a version designed for educators in low-bandwidth regions.

KNOWN LIMITS Until the low-bandwidth version exists, LUDIA requires a reasonable connection and does not serve offline settings. The 13 languages widen reach but do not by themselves make LUDIA contextually fit: a machine translation carries the assumptions of the English it came from, so an educator working in any of the other languages may receive guidance shaped for somewhere else. UDL originated in the United States, and its uptake in national policy is uneven. How useful LUDIA proves in a given curriculum or education system depends on the educator adapting it into their own setting, and that is work we are asking of them rather than doing for them.

Inclusive access

DESIGN COMMITMENT LUDIA is built to WCAG 2.2 Level AA and to the accessibility requirements of the European Accessibility Act, Directive (EU) 2019/882, whose technical requirements for ICT are set out in EN 301 549 (ETSI, 2021), and we publish an accessibility statement describing what that means in practice: keyboard operation throughout, interface roles and states exposed using WAI-ARIA (W3C, 2023a), screen reader support tested with NVDA, Orca, VoiceOver, and TalkBack, contrast meeting or exceeding 4.5:1, browser zoom to 400 percent with content reflowing alongside a text size setting, six notification placements, adjustable dismiss timing, light and dark schemes, and Lexend Deca as the body font, chosen because its letter shapes are open and easy to distinguish. The absence of an account removes a barrier that is easy to overlook: no email address, no institutional login, no approval from an administrator, and no record of who asked what. Educators without institutional backing use LUDIA on exactly the same terms as anyone else, whether they work in a school, a community or non-formal programme, a refugee or emergency setting, or on their own.

NOT YET IN PLACE Accessibility is tested but not certified. We test with automated tools, manual keyboard-only passes, four screen readers, contrast checks against WCAG 2.2 Level AA ratios, and informal feedback from educators who use assistive technology. That feedback is not structured testing with disabled educators, which we have not done and which Part Four sets out as work still to come. There is no independent accessibility audit yet. An Accessibility Conformance Report against WCAG 2.2 Level AA, following the VPAT 2.5 template, is being completed by a certified accessibility specialist we have engaged. That report is self-authored, as conformance reports are: it names its evaluator and records the testing carried out, and it does not stand in for an audit. An independent audit by a partner with no part in building or reviewing LUDIA is planned. When it happens we will name who carried it out and publish what they found, including anything they find against us. Nobody has assessed the accessibility of what LUDIA produces, as distinct from the interface presenting it: the reading level of a reply, its length, how it is structured for someone reading by screen reader. We treat that as an accessibility gap, not a question of tone, and the plain language version described later is the first part of our answer to it. What has not been done at all is an audit of LUDIA's outputs for cultural or linguistic bias. That risk is concrete: the model's examples and assumptions may skew towards well-resourced, English-medium, Global North schooling.

How LUDIA reaches a learner

LUDIA has no contact with children, and it still reaches them. It reaches them the way any professional resource does, through the design decisions an educator makes after using it. That pathway is not incidental. It is what LUDIA is for, so we should say what travels along it. Two things named elsewhere in this document travel it directly: the model's suggestions may be calibrated for well-resourced, English-medium, Global North schooling, and the model may agree with a design an educator brings instead of questioning it. Neither has been tested. What stands between those risks and a learner is the educator's own knowledge of that learner, which is real and which the whole design relies on, and the fact that LUDIA offers options to weigh, not instructions to follow. What does not stand between them is anything we have built or measured.

We have completed a Data Protection Impact Assessment, and that is an assessment of data, not of children. A structured review of LUDIA against General Comment No. 25 on children's rights in relation to the digital environment is work we have not done. It belongs in Part Four, not in a sentence here.

A tension inside the architecture

There is a tension in all of this, and this is the place to say it. Because LUDIA holds nothing, it cannot receive a safeguarding disclosure. There is no account, no record of a session, and no way for us to know who wrote anything or to act on it afterwards. If an educator types something indicating that a child is at risk, nothing about it reaches us, and by design nothing can. The structure that protects every user is the same structure that makes that disclosure impossible to act on, and the two cannot be separated.

We are not going to resolve it by weakening the architecture, because a tool that could act on such a disclosure would be a tool that retains and identifies. What we can do is be exact about what LUDIA is not. It is not a safeguarding route and must never be treated as one. The route is the educator's own institutional process, and it always was. Designing LUDIA to say so at the moment it matters, instead of carrying on with the design conversation as though nothing had been said, is work in front of us, and it is set out in Part Four. This is a tension we have named and are working through, not a problem we think we have solved.

Part Two: The European Commission's guiding questions

The guide. Guidelines on the ethical use of artificial intelligence and data in teaching and learning for educators, revised edition. European Commission, Directorate-General for Education, Youth, Sport and Culture. Publications Office of the European Union, Luxembourg, 2026. <https://doi.org/10.2766/7967834>. Licensed CC BY 4.0. Eight requirement areas, addressed to teachers and school leaders, situated within the EU Artificial Intelligence Act and the GDPR.

The Commission's guidelines differ in kind from the other three, in two ways that change how we can respond to them. First, they are addressed to teachers and school leaders, not to evaluators or providers. The questions belong to educators and learning leaders, and it would be a category error for us to answer them on their behalf. So this section does not score anything. It gives educators what they would need to answer the teacher-facing questions themselves, and it names the ones that belong to a school and not to us. Second, these guidelines sit alongside binding law. They are non-binding themselves, and they state explicitly that they do not provide guidance on the enforcement of the AI Act. But they set out the legal terrain, and any tool entering a European classroom should be able to say where it stands on it.

Where LUDIA sits in the AI Act

One question sits underneath everything else here, because every obligation in the Act attaches to a role. LUDIA is placed on the market under our own name, built on a general-purpose model supplied by others. On our reading that makes us the provider of an AI system, not the deployer of someone else's, and it would make the transparency duties described below ours and not a school's. We have not taken legal advice on that question and are not currently in a position to. We would rather write it down than leave a reader to work out which role our reasoning quietly assumes. Everything that follows should be read with that unsettled.

Here is our reading. Treat it as a starting point for a school's own due diligence, not a legal determination. The Act's education-related prohibition concerns systems that infer emotions in educational settings, and LUDIA performs no emotion recognition of any kind. The high-risk category for education under Annex III covers systems used to determine admission, to evaluate learning outcomes, to assess the level of education a person will receive, or to monitor prohibited behaviour during tests. LUDIA does none of these. It does not see learners, does not process learner data, produces no scores, predictions, or rankings, and makes no decision that reaches a learner except by way of an educator's own judgement. LUDIA is a design thought partner for adult professionals. On our reading it therefore falls outside both Article 5 and Annex III, and its obligations are transparency obligations: users

must be told they are interacting with an AI system, and generated content must be identifiable as such. LUDIA identifies itself as an AI system at the point of first interaction, and generated text is marked in a machine-readable format where the delivery channel carries such marking, as Article 50(2) requires. That is a conditional, and a conditional attached to a compliance statement should be sized rather than left to stand: an inventory of which channels carry the marking and which do not will be published with the next revision of this document. We treat all of this as a floor, not a courtesy. Two further points matter here. The obligations attaching to Annex III systems apply from 2 August 2026, and further Commission guidance on high-risk practices is expected. A school satisfying itself about any AI tool should check where things stand now, not rely on ours.

Where LUDIA sits in the GDPR

LUDIA sets no cookies and requires no account. Chat content is held in the browser's session storage for the length of the session and cleared when the tab closes; it is never written to LUDIA's servers. Preferences such as theme and text size sit in local storage on the device, as does a simple record that the user confirmed they are 18 or older. One cookieless analytics tool, Vercel Web Analytics, counts aggregate page views using a short-lived hash generated from the request, discarded after 24 hours rather than kept as a persistent identifier, and it receives no chat content. Country-level figures are available without the underlying IP address being stored or tied to them. Because nothing is stored on or read from the device for tracking, no cookie banner is needed.

That substantially lowers what a school has to consider without eliminating it, and the residual point deserves stating plainly. Whatever an educator or instructional designer types, and any file or image they attach, is transmitted to the model provider so that a reply can be generated. An uploaded file is never stored by LUDIA in any form at any point: there is no file storage service, nothing is written to disk even temporarily, and the file content travels inside the message itself, so no copy exists afterwards to be deleted. One nuance matters here, because nothing is held on our side: the conversation is present in full with each new message, so an attached file travels with it until it ages out of a size limit and is replaced with a short note. The provider, which holds SOC 2 and ISO 27001 certification, processes the request in working memory and deletes it once the reply is returned. Retention is determined by how the host institution configures its Google Cloud project. Google does not use prompts or responses to train or fine-tune its models, and this deployment is not subject to prompt logging for abuse monitoring. If a learner's name, a photograph of their work, or a profile document is entered, that becomes personal data leaving the school through a route the school has not assessed. Nothing in LUDIA's architecture prevents that. Our privacy position is straightforward: "Privacy is a human right and our shared responsibility. Share the question, not the student. If the words or images shared would let someone identify a real child, remove them before pressing send."

Article 28 of the GDPR requires a written contract between a controller and a processor before the processor handles any personal data. That contract is in place: Google's Cloud Data Processing Addendum governs the processing carried out on Vertex AI, and the commitment that nothing is used for model training is part of those terms. The same

question applies to the two other services that see anything at all: Vercel, which hosts LUDIA and provides the cookieless analytics, and Bitly, which counts clicks on the short links used on our resource pages.

Human agency and oversight

LUDIA's educational function is design support. It helps educators and learning leaders think through the design of a learning experience, a resource, or an environment. It recommends nothing to a learner, assesses nothing, and predicts nothing. Every output passes through an educator before it can affect anyone, which means the oversight this section asks about is not a feature we added but the only way the tool can operate at all. That is a weaker safeguard than it sounds, and the reason is automation bias: a person reviewing the output of a tool they believe to be competent catches less than a person thinking without one. The educator most likely to accept a poor UDL suggestion is the educator who trusts that LUDIA knows UDL. We have built nothing against that beyond framing every response as an option to weigh, and we name it here because a human in the loop is not self-evidently a control. Two of this section's questions belong to a school, not to us: whether teachers have had professional development on using AI-supported tools, and who is responsible for reviewing AI outputs.

Transparency and explainability

It is worth being plain about how much the word "engineered" should carry. LUDIA is a general-purpose model working within design decisions we made. We have not trained a model of our own and we do not claim to have. LUDIA draws on the UDL framework to shape examples and suggestions for the situation described, and it directs educators to CAST's published guidelines, which are the authoritative source.

Transparency also covers what a tool is, and not only how it works. LUDIA identifies itself as an AI system before the first exchange, takes the pronoun it, and is drawn as a chatbot rather than a person, for the reasons set out above.

LUDIA's reasoning is explainable in a way that many AI tools' reasoning is not, because its suggestions rest on a published framework and an educator can judge for themselves whether the reasoning holds. We can also name the model. Currently, LUDIA runs on Gemma, Google's openly published model, operated on Google Vertex AI. Open weights can be inspected in a way a closed API cannot, but open weights are not open training data: they tell an examiner more about how the model works than about what shaped it, which is the part that bears on bias. We can change the model, and the code change would be small; the real cost of changing it would be revalidating that LUDIA still behaves as intended, because the instructions are tuned to how this model responds. There is a further thing transparency should mean here, which this document has not offered until now. LUDIA's behaviour comes from a general-purpose model working inside instructions we wrote. This paper says in three places that its claims can be verified by using LUDIA, and that claim is only as good as what we are willing to say about how LUDIA is shaped. So we are publishing a behavioural specification: a general account of how LUDIA is instructed to behave, including which model it currently runs on. It carries a version number and a date. A change to the model or to the

underlying instructions produces a new version of it, and the change is recorded there. That specification is also where the process for changing the model lives, which closes a gap named in Part One. What no specification can offer is an account of why the model produced a particular sentence, or a reproduction of the underlying code and instructions themselves. LUDIA's output is best treated as a colleague's suggestion to be weighed, which is also how responses are framed at the point of use.

Diversity and inclusion

This section addresses accessibility, universal design, and stakeholder participation, and it is the closest the guidelines come to LUDIA's own subject matter. LUDIA is designed to WCAG 2.2 Level AA with a published accessibility statement, and it exists to help educators apply universal design, not simply to embody it. The questions about learners without access outside school do not apply directly, since learners do not use LUDIA. They apply in translated form to educators: until the low-bandwidth version is developed and shared, an educator without a reliable connection is an educator LUDIA does not reach. On participation, beta testing by volunteer educators across languages and contexts is the honest answer for this relaunch.

Fairness and non-discrimination

Nobody has tested LUDIA for bias, across learner groups or educator groups. That gap belongs at the front of this section, not buried in it. The risk is real and specific. The model LUDIA is built on may produce examples, assumptions, and defaults skewed towards well-resourced, English-medium, Global North schooling, so an educator working outside those conditions may receive guidance quietly calibrated for someone else's classroom. Multilingual delivery widens reach but does not by itself address this, and machine translation introduces a second layer of risk, which we disclose instead of burying. Our present mitigations are design-level: suggestions grounded in an explicit framework, sources given, and everything framed as an option to weigh. Until testing exists, an educator's own knowledge of their learners remains essential, alongside what UDL itself contributes: a way of identifying barriers and building awareness of them that does not depend on LUDIA alone.

Societal and environmental wellbeing

LUDIA carries no engagement mechanics: no streaks and no badges. The in-app notices that exist are functional rather than promotional, and their placement and the dismiss timing are set by the educator. LUDIA is meant to be closed once the work is done. Since learners do not use it, this section's questions about learner motivation, mood, and dependence do not apply directly, though the underlying concern translates upward to professional over-reliance, which we take seriously and address under educational impact above. On the environmental question, which this section asks and few tools answer, our position is set out in the sustainability section below.

Privacy and data governance

Covered above. The version an educator can act on: LUDIA collects nothing about educators or their learners, retains nothing once the tab closes, and the only data leaving the browser is what someone chooses to send, whether typed or attached as a file or image. Short links on our resource pages run through Bitly, which counts clicks but does not identify anyone. The clause about what is typed is where the responsibility sits, and it sits with the person typing.

Technical robustness and safety

There is more here than there was during LUDIA's first three years. TLS 1.2 or higher protects data in transit, no user data is written to durable server-side storage, access is role-based with quarterly reviews and multi-factor authentication, and there is centralised audit logging and a documented incident-response plan. Separate staging and production environments are in place, with automated checks before anything reaches production. Abuse and runaway cost are bounded by three layers: bot detection on every message before it reaches the model, network-level rate limiting at twenty messages a minute and three hundred an hour from any one source, and a hard ceiling on model spend of one hundred dollars a month. Monitoring runs on the hosting platform and, independently of it, an external service checks the live site every few minutes, so a full outage would be caught even if the platform itself failed. There is a public status page at status.askludia.com. Accessibility and fault reports go to accessibility@askludia.com. They are acknowledged within two business days and answered with concrete next steps within five business days, with an alternative route to the content offered while a fix is in progress.

What is still missing is external scrutiny. Nobody outside has audited LUDIA for reliability, and nobody will have by relaunch. We have published no uptime commitment. There is also a ceiling we should name rather than discover in public: the model provider caps concurrent requests to a model at two hundred by default. Each message triggers two of them. So errors would begin somewhere around a hundred educators messaging at the same instant. What those educators would see today is a generic failure message rather than one explaining what has happened and when to try again, which is itself worth fixing. That is a capacity limit, not a scaling plan, and it is one of the things the relaunch will tell us about. The stakes are low: LUDIA runs in a browser, stores nothing, and touches no school system, so if it goes down, nothing is lost. That is a benefit of the design, and no substitute for a procedure.

Accountability

Two named people answer for LUDIA. The arrangement has an obvious strength and an obvious limit. Nobody is hiding behind a company, so if LUDIA gives poor guidance, we are the ones who answer for it, and we can be reached directly at hello@askludia.com. The limit is capacity. We cannot offer a support desk, escalation tiers, or anything a procurement office could take to a lawyer. The one commitment we do make is the response time for accessibility and fault reports set out above, which two people can hold and which we will report against rather than quietly let slip. Schools should weigh that. No amount of good intent closes that gap for a school that needs a contract.

On the five ethical considerations

Beneath the guiding questions sit five ethical considerations: human dignity, fairness, trust, academic integrity, and justified choice. The last is the one we would draw attention to, because it names something the other guides leave implicit. Justified choice asks that decisions about AI in a school be transparent, participatory, and explainable. It puts the weight on how a decision was reached, not on which tool was chosen, and it is, we think, the strongest available answer to what Evidence of Trust could really mean: not a property certified in a tool, but a decision an educational community can account for. What an educator needs in order to justify a choice is a clear account of what a tool is, what it is not, and what is not yet known about it. That much we can offer today.

Part Three: The EdTech Quality Indicators

The guide. EdTech Quality Indicators Guide. EdTech Quality Collaborative, June 2026. Licensed CC BY-NC-ND 4.0. Developed by 1EdTech, CAST, CoSN, Digital Promise, ISTE+ASCD, and SETDA. Five indicators: Safe, Evidence-Based, Inclusive, Usable, and Interoperable.

Everything in this section is a design commitment. Almost all of it can be verified by using LUDIA and observing what it does.

Safe: users' privacy is protected

A safe tool "adheres to industry standards, data minimization, and data transparency to protect user privacy and prevent data breaches" (EdTech Quality Collaborative, 2026, p. 2). LUDIA takes this about as far as a tool can. LUDIA requires no account, sets no cookies, and does not write chat content to our servers, nor collect data for model training. The conversation is held in the educator's own browser for the length of the session and is cleared when the tab closes, and LUDIA keeps no memory of it across sessions. An exchange worth keeping can be downloaded and uploaded again later to continue where it left off.

There is something this indicator does not cover. It appears in our summary table, so it belongs here too. What an educator types is transmitted to the model provider so that a reply can be generated. Retention is determined by how the host institution configures its Google Cloud project, Google does not use prompts or responses to train or fine-tune its models, and this deployment is not subject to prompt logging for abuse monitoring. The contract position on that is set out under Privacy and data governance, and it is settled.

Evidence-based: grounded in research

A tool that meets this indicator is "grounded in rigorous research, educational standards, and empirical validation" (EdTech Quality Collaborative, 2026, p. 2). This indicator requires the most careful reading of the five, because it can be heard as a claim we are not making. LUDIA is not itself a research tool. LUDIA neither generates nor evaluates evidence, and it has not been empirically validated. No study of LUDIA's effects exists. What it does is connect educators and learning leaders with an evidence-based approach: Universal Design for Learning, CAST's research-based framework for reducing barriers and centring learner agency, now in its third version of the Guidelines, together with the asset-based pedagogies and practices that align closely with UDL. The research base sits underneath what LUDIA points to, not underneath LUDIA itself.

A note on UDL's own evidence base

UDL's evidence base is contested, and a document about evidence ought to say so before it is asked.

Boysen (2024) reviewed the empirical studies CAST cites in support of three Guidelines and reported that most did not offer learners a choice, most did not measure learning, and none concerned brain function. A follow-up review found that around one percent of the sources CAST cites represent neuroscience, and that three quarters of the Guidelines had no neuroscientific source at all (Boysen et al., 2025). A related line of criticism holds that multiple means can be read as a version of learning styles, an idea that does not survive testing. None of this is frivolous, and we will not pretend otherwise.

The criticism does not fall evenly across UDL. The neuroscience framing, the three networks of the learning brain, takes the most damage, and it is the layer LUDIA leans on least. CAST itself calls it a simplified model. Less exposed is the design proposition: that learner variability is predictable, that inflexible design creates barriers, and that anticipating a barrier beats retrofitting once a learner has already met it. It is that proposition the CRPD Committee endorsed in General Comment No. 4, when it directed States Parties to the Convention, ratified by more than 190 states, to adopt the universal design for learning approach. An interpretation of a human rights obligation is not an empirical finding, and it does not answer the question Boysen asks. It does mean the case for the design proposition was argued on access and equality. It never rested on a study of choice or of brain imaging.

The field is not ignoring the question. Boysen's paper drew a published reply from Edyburn (2024), one of UDL's longest-standing scholars. Basham and Lowrey (2025) call for a global community using scientifically rigorous methods to identify universal foundations for learning. CAST's own review for Guidelines 3.0 found the existing research base narrow across most of the dimensions examined and set out to widen it. A field that argues about its own foundations in public is in better standing than one that does not.

We co-created LUDIA, and we are not going to hold UDL to a lower standard than we hold LUDIA. LUDIA connects educators to a published framework for identifying and reducing barriers, one with legal standing, forty years of practice behind it, and a research base under active and legitimate challenge. LUDIA does not claim that its suggestions have a direct impact on learning gains. Neither do we, and we co-created it.

Inclusive: built to remove barriers

An inclusive tool "prioritizes accessibility and equitable design ... mitigating algorithmic bias and stereotypes" (EdTech Quality Collaborative, 2026, p. 2). The responses point educators towards equitable, inclusive practices that "proactively remove barriers to learning" (EdTech Quality Collaborative, 2026, p. 12). LUDIA is designed to comply with WCAG 2.2 Level AA and publishes an accessibility statement describing what that means in practice, because professionals in education have the right to inclusive tools and spaces, digital and face-to-face (CRPD, 2006, Articles 9 and 27).

A note on the language we use: these documents move between identity-first phrasing, as in disabled educators, and person-first phrasing, as in children experiencing disability, because preference differs between communities, languages, and regions, and because the instruments we cite differ too. Where we quote a source we keep its wording. Elsewhere we

follow the convention most used by the people being described, and we will change it when we are told we have it wrong. Building to a standard is one thing. Testing with the people the standard protects is another, and that work is still ahead of us.

Usable: easy to work with

A usable tool "focuses on intuitive design ... to eliminate technical barriers for teachers and students" (EdTech Quality Collaborative, 2026, p. 2). LUDIA runs in a browser with nothing to install and nothing to sign up for, in 13 languages, and each exchange includes inline links to relevant resources and next steps. What we cannot report is whether any of that makes LUDIA easy to work with, because we have carried out no usability testing we could offer as evidence. Volunteer educators used the tool before relaunch and their reactions shaped it, and none of that was recorded in a form anyone else could check. This indicator is therefore unanswered, and answering it is among the first things we want to find out.

Interoperable: an indicator written for a different kind of tool

This indicator is about tools that "utilize open data standards to seamlessly connect with existing [institutional] technologies and securely aggregate instructional data" (EdTech Quality Collaborative, 2026, p. 2), the systems that pass student rosters and grades from one education platform to another (EdTech Quality Collaborative, 2026, p. 13). LUDIA does not integrate with the platforms a school, college, or university already runs, and it does not aggregate or pass along instructional data. Because it holds no data of any kind, there is nothing to connect and nothing moving between systems. It will not, either. We should say that flatly. There is a way of writing this section that turns the gap into a virtue, and it would not be true. This indicator presumes a tool that holds data and sits inside an institution's systems. LUDIA holds nothing and sits inside nobody's systems, so there is nothing to connect and nothing to pass along, and no amount of development will change that without changing what LUDIA is.

What a school gives up is real and should be named: no single sign-on, no roster sync, no gradebook connection, and no instructional data flowing into the systems leaders read. An institution that needs a tool inside its existing stack should not choose LUDIA. What it gets instead is a tool that works from the first click, with nothing to procure, configure, administer, or write a data-sharing agreement about, and that cannot leak what it does not hold. That is a trade, not a score. Of the indicators in these four guides, this is the one written for a different kind of tool. Saying so is not the same as claiming a pass, and it is more accurate than recording a failure that misdescribes the design.

Part Four: What we will count as evidence

The guide. Needs-Based EdTech Evidence Mapping. Lindroos Cermakova, A., Havinga, B., & Clary, A. European EdTech Alliance, Bielefeld, March 2026. Licensed CC BY-NC-SA 4.0. Five evidence domains, sorted by the purpose evidence serves rather than by methodological strength.

The Needs-Based Evidence Mapping sorts evidence by the purpose it serves, not by how impressive the method looks. Its most useful feature is that it shows where a portfolio is thin. One distinction has to be drawn first, because this is exactly the sort of place where the report warns that evidence changes meaning as it travels.

The LUDIA that ran on Poe between 2023 and 2026 accumulated a body of independent attention. Researchers and practitioners have referenced it in peer-reviewed journal articles, a practice brief, conference proceedings, an industry report, a professional journal, a white paper, blog reviews, videos, and a substantial number of podcast episodes, none of them commissioned or paid for by us. We keep that collection public and citable under "Learn About LUDIA" at askludia.com. But we want to be exact about what it is. It shows that educators and researchers found LUDIA worth writing about, recording, discussing, and demonstrating to colleagues. What it does not show is that LUDIA works. None of it measured that. A citation is not a finding, and reach is not the same as impact.

So here is what we intend to gather from August 2026, domain by domain. We are setting it down now so that it can be checked against what we actually produce. This is an order of priority, and it now carries a date. We will publish a revised account of this document in August 2027. By then we intend the following to exist: in domain 2, a documented account of the beta and the first results from the sycophancy testing described below; in domain 3, the written safeguarding and incident-response procedures and the published behavioural specification; in domain 4, the Accessibility Conformance Report and an independent audit arranged; in domain 5, a plain account of what LUDIA costs to run. Domain 1 will be unchanged, for the reasons given there. Where any of it has not happened, that revision will say so and say why, which is why it is written here. We are two educators. Gathering evidence well takes time and expertise we will not always have to hand, where that is true we say so, and where we fall short the next revision of this paper will show it.

1. Learning Impact & Effectiveness

NOW Nothing. No trial, no pre-post measure, no attainment data, and no proprietary study standing in for one. This is the domain in which claims are most often stretched in our sector, and we make none.

INTENDED Little, and cautiously. LUDIA is one input into an educator's design process, and any effect on learners is mediated by that educator, their learners, and their setting. We do not think a credible attribution of learner outcomes to LUDIA is available to us, and we will

say so permanently instead of promising a study we cannot run. If this domain fills, it will be through independent research rather than through anything we claim, and we are working to make that research possible. We are seeking partnerships with university researchers who would design and run such a study on their own terms. Nothing is contracted, and we will name any partner here once it is. That position has two consequences, and both belong on the record. Any independent researcher who wants to measure LUDIA's effects will have our cooperation, whatever access we can give without breaking the privacy architecture, and no right of review over what they publish; we will link the result from this document whether it flatters LUDIA or not. In the meantime a learning community is left with the justified choice standard described in Part Two. Not a claim that the tool works. An account clear enough about what it is, what it is not, and what is unknown about it that a decision can be made and defended. That is a lower bar than proof of impact. It is the one we can actually meet, and naming it is better than letting a reader assume the higher one has been cleared.

2. Experience, Usability & Pedagogical Fit

NOW Beta testing, undocumented so far. Volunteer educators have been testing the relaunched LUDIA in their own planning and design work, across languages, education systems, and cultural contexts, and what they found shaped the build. None of it is documented yet, so it cannot yet be offered as evidence to anyone but us. Documenting and sharing what we learned from it is part of our plan, not an afterthought, and it is the first thing we intend to publish below.

INTENDED Documenting and sharing what the beta testing already found comes first: who took part, across which languages and contexts, what they told us, what changed in the rebuild as a result, and what we chose not to change. That account is part of our plan, not a possibility, and closing it is the first commitment listed in Part Four's timeline below. Beyond that, this is where we expect the portfolio to build most substantially going forward, and it is the evidence the report describes as valued by educators and undervalued in formal evaluation. We want to know which questions educators actually bring to LUDIA, where its responses are useful and where they miss, whether the inline links to relevant resources and next steps are followed or ignored, and whether the tool fits into the moment of planning or interrupts it. Above all we want to know whether the Four T's actually happen, and whether educators reach Transfer. Telling, tinkering, and tweaking leave traces within an exchange. Transfer, as described above, happens after the exchange closes and leaves no trace in the tool at all, which is precisely why it is the hardest of the four to gather evidence about. It is also the one the whole design is aimed at. An educator or instructional designer who takes the first response and goes is using LUDIA as a solution engine, which is the one thing it was built not to be. Other people have walked colleagues through that routine with the old LUDIA, which suggests it travels. Whether it is common or rare with the new version of LUDIA is an important question we plan to ask.

One thing here we will look for rather than wait to notice. Sycophancy, named earlier in this document, is the failure that would empty LUDIA out. A tool that agrees with the design an educator brings it has removed the one thing it was built to add. We will put designs with known barriers to LUDIA, in each of the languages it offers, and record whether it names the

barrier or agrees with the plan. Some of those runs will state a preference first, and some will defend the flawed choice rather than simply propose it, because that is the harder case and the more common one. We will report what comes back whichever way it goes. A poor result is a finding about LUDIA, not an argument for a different test.

3. Safety, Ethics & Data Governance

NOW Built in, but not written down. No accounts, no cookies, no chat content held on our servers, no data collected for model training, and no memory across sessions. Very little remains to be governed, and a claim of this kind is verifiable by inspection in a way that a policy document is not.

INTENDED The documented layer this domain also expects. Written safeguarding, incident-response, and model-change procedures, a clear public statement of what leaves the browser and where it goes, and a structured review against General Comment No. 25 covering the pathway from a suggestion to a learner. That is drafting work, not discovery work, and it does not depend on the relaunch.

4. Equity, Accessibility & Inclusion

NOW Design evidence and a published statement. WCAG 2.2 Level AA as a design target, with a published accessibility statement setting out what was built to meet it. No login and no cost as barriers are removed. Thirteen languages. UDL as the substantive commitment underneath. Building to a standard and describing it is not testing. We have not tested with disabled educators.

INTENDED An independent accessibility audit, commissioned from a partner with no role in building LUDIA, accessibility testing with disabled educators, and a structured review of LUDIA's outputs for cultural and linguistic skew, including with educators working outside English-medium and well-resourced settings and in the machine-translated languages. Where the model proves poorly calibrated for a context, we want it written down and not quietly absorbed.

5. Scalability, Implementation & Sustainability

NOW A favourable architecture and an unquantified dependency. A tool with nothing to install, nothing to integrate, and nothing to store scales without the usual costs. We could point here to the earlier version's reported reach across more than 120 countries, and we are not going to. That figure came from a third-party page's traffic data we cannot confirm, and it describes a tool that no longer exists in that form. It is not evidence about this one, and this is exactly the section where a report about how evidence changes meaning as it travels would expect us to use it anyway. Against that, LUDIA rests on a third-party model whose behaviour, pricing, and availability we do not control. We can now put a rough figure on part of that, from our own costing against the provider's published rates in July 2026. A short conversation costs around eight tenths of a cent, most of which is the core instructions being resent with every message rather than the conversation itself, which puts a thousand conversations at roughly eight dollars a month and ten thousand at roughly eighty. Those are estimates from published

pricing and not measurements. Conversations vary, an attached image costs noticeably more, and languages needing more tokens to carry the same meaning cost more again.

INTENDED A plain account of what LUDIA costs to run and what happens if that changes, and a stated position on continuity: what we would do, and what we would tell people, if the model changed underneath us or the tool had to stop. A tool designed to public-good terms that disappears without notice is not honouring those it is built to serve.

One thing, well

Designers are never finished. That is not a complaint. That is the condition of the work, and it is why these four guides are a beginning and not a verdict.

There are two ways to keep working on a tool, and they lead to different places. One is to add. More features, more integrations, more tiers, more reasons to stay. That is the typical path in EdTech, and it is how a tool offered at no cost becomes a paid one. The other is to leave the tool doing the one thing it does, and to spend the effort on the distance between the tool and the people who cannot yet reach it.

LUDIA takes the second path. LUDIA does one thing, and keeping it to one thing is a design commitment rather than a stage it will grow out of. LUDIA connects educators with the UDL Guidelines and the asset-based pedagogies that work alongside them, one situation at a time. It will not become a marking assistant, a data dashboard, or a suite. Every capability we decline to add is a capability we do not then have to secure, price, or maintain, and it is attention we can spend somewhere more useful. As the design evolves, the work is not adding capability. The work is reducing barriers, and the four guides are how we find them.

Two barriers are already named. The first is an accessibility problem in what LUDIA produces, and it should be called that, not treated as a matter of style. Conformance to WCAG 2.2 Level AA covers the interface an educator moves through. It says almost nothing about the replies that interface delivers, and the replies are the substance of the tool. LUDIA's language can run academic, and its responses can run long. An educator can ask LUDIA to simplify a response, but asking is itself the barrier: it puts the work of reducing the barrier back on the person meeting it, which is exactly what LUDIA exists to avoid. We are building the plain language version described below because a workaround is not a fix. A plain language version built around the UDL Access Guidelines is part of our development plan, and we treat it as accessibility work, not as a way of widening reach. LUDIA requires an internet connection, which excludes anyone working where connectivity is scarce or costly, and a low-bandwidth version answers that. Neither adds anything LUDIA can do. Each removes something standing between LUDIA and an educator or instructional designer who cannot reach it yet. More languages do the same work, and so does asking for no account, no cost, and nobody's permission.

Voice is the same principle running the other way, and it is the place where we have not yet found an answer we are willing to defend. LUDIA has no speech-to-text and no text-to-

speech. Most voice features in chatbots route the audio, or the chat text, to a separate service run by another company, which may keep it, log it, or train on it, sometimes on servers outside Europe. Adding that would break the promise that LUDIA saves nothing and builds no profile, so we did not add it. What we do instead is point to the voice tools already on the educator's own device, which run locally: Voice Access on Windows, Voice Control and Dictation on Mac, Dictation on Chromebook and iPad, Voice Typing on Android.

That arrangement is unsatisfactory and we would rather say so than let it pass as a principled decline. Built-in voice is not a WCAG 2.2 requirement, so this is not a conformance gap. It is a design gap, which by our own standard is worse. An educator who works by voice, whether because of a dexterity or motor disability or for any other reason, arrives at LUDIA and is told to go and set something up on their own machine before the tool is easy to use. The effort has landed on the person instead of on the design. That is the move UDL exists to refuse, and pointing at the operating system does not transfer the responsibility for it. Ease of use belongs at the front, in the thing we built.

So we hold both facts. The privacy reason for not routing audio to a third party is real and we are not going to abandon it. The current position is still an obstacle, and it is ours. We are working on a way to do this that keeps the audio on the device and sends nothing to anyone, and until we have one this stands as a known barrier named plainly in the text. It sits alongside the plain language version and the low-bandwidth version, all three still outstanding. That is what growth means here. Growth means a shorter list of people LUDIA cannot reach. The feature list stays where it is. UDL asks the same thing of any design. The barrier sits in the design, not in the person who meets it.

Designing for sustainability

Beyond the five pillars, the eight requirement areas, and the five domains, LUDIA is designed to use less. UNESCO's Recommendation on the Ethics of Artificial Intelligence (UNESCO, 2021a) names "environment and ecosystem flourishing" as a core value. The EdTech for Good Framework asks directly about the energy, compute, hosting, and storage implications of deploying a tool. The European Commission's guidelines ask educators whether they have been told anything about a tool's environmental impact at all.

There is a prior question here that we should not skip. The largest energy decision was using a generative model at all, when a well-made static resource costs almost nothing to serve. Our answer is that the UDL Guidelines already exist as a static resource, and the barrier we set out to reduce was not access to the text but the distance between the text and a particular design problem. That is a judgement, not a finding. Having made it, we try to spend as little as the work requires, and to be accurate about where the spending goes.

Nothing accumulates between conversations. LUDIA saves no past chats and carries nothing from one session into the next, so a returning educator starts from a clean slate rather than from a growing history. That is a privacy decision and an energy decision at once. Within a

conversation, the design pushes towards a purpose and an ending rather than an open thread: one situation at a time. Generating a long response costs more energy than generating a short one, with reductions in the range of 25 to 60 percent reported for shorter outputs that remain useful (Poddar et al., 2025).

There is a cost we do not currently control, and it would be misleading to leave it out. The instructions that shape how LUDIA behaves travel with every message, because the provider offers no way to cache them. We would reduce that if the option existed, and we will take it when it does. So the honest claim is narrower than a claim to efficiency in the abstract: LUDIA is built for one kind of work and does not accumulate. Consider a short trip across town. A truck will do it, and a truck can carry anything. An e-bike is built for the trip. LUDIA is built like the e-bike, not because every mile is cheaper but because it is not carrying what it does not need. That is an argument about design, and it is not a measurement. We have not measured LUDIA's energy or water use, we hold no figure for the cost of a single exchange, and until we do, everything in this section is reasoning about how the tool is built rather than evidence about what it consumes. Producing that figure is in Part Four, domain 5.

This follows UNESCO's call for AI that uses less and works for everyone, including teachers and learners in places with fewer resources, set out with University College London in *Smarter, Smaller, Stronger* (UNESCO & University College London, 2025) and in the *AI for Environment and Ecosystems Toolkit* (UNESCO, 2026). LUDIA is not part of UNESCO, and UNESCO has not approved or endorsed it. We cite these sources because they shape how we build.

What happens next

This document is a starting position, and we would like it to age visibly. The commitment is ongoing rather than delivered. We would rather state what we intend and be seen to fall short of it than promise a research programme we cannot staff.

In the months after relaunch we intend to close the gaps named above: a documented safeguarding and incident-response procedure behind the reporting address we already publish, a stated process for changing the model LUDIA runs on, and a public account of the beta testing that shaped the rebuild, including what we chose not to change. We intend to build the plain language version and the low-bandwidth version, because both take a barrier away. And we will begin gathering in the domains set out in Part Four, starting with experience, usability, and pedagogical fit, before publishing a revised version of this document that reports what we found, including the parts that do not flatter the design. Where a guide's question has moved from unanswered to answered, we will say so. Where it has stayed unanswered for another year, we will say that too. What will not change is the shape of the tool. If LUDIA is doing more in three years, we would rather it were doing the same thing for more people.

A closing note on trust

We do not regard LUDIA, or this design and evidence statement, as a finished achievement. A badge is earned once and displayed forever. Earning trust is more nuanced and more demanding, held by particular people in particular settings and open to revision whenever either of those changes. The learners and professionals affected by EdTech tools depend on the judgement of the people designing those tools. In designing LUDIA, we therefore treat human rights as imperatives rather than add-ons: the right to inclusive education, the right to accessible tools and information, and the right to privacy.

We are publishing this as an invitation to others building for education. Education is more than a marketplace, and EdTech is more than an industry. Those three rights are why LUDIA exists, and Universal Access to Universal Design for Learning is the goal we set ourselves. A priority is only a priority if it can be found in the thing that was built, and ours are in the architecture: no cost, no account, no cookies, nothing stored, WCAG 2.2 Level AA, 13 languages. The terms are workable, and that is the part worth reporting.

Trust is earned continuously, from learning leaders and from the members of the learning communities they serve. We are grateful to the people who wrote the four guides, and we will keep working from them. That all four are published at no cost matters twice: a developer without a budget can still read them, and anyone can hold us to them.

CAST works under the motto "Until learning has no limits." That is what we are designing towards.

Funding and competing interests

LUDIA received no external funding. It has no investors, no advertisers, no sponsors, and no commercial relationship with curriculum vendors, and it generates no revenue. Development and running costs are met by the authors. The authors created LUDIA and have a non-financial interest in how it is received. In 2025 the authors received CAST's Anne Meyer UDL Design Award; the award carried no funding, and CAST has neither reviewed nor endorsed this statement. The award also creates an overlap worth naming here before someone else finds it. CAST is a named co-developer of the EdTech Quality Indicators Guide, one of the four guides this document reads LUDIA against. Part Three gives our reason for choosing it. The choice was still made by people that organisation had given an award, and anyone weighing that section should know it. None of the authors of the four guides read in this document has reviewed or endorsed LUDIA.

Citation and references

SUGGESTED CITATION Stark, B., & Rostan, J. (2026). LUDIA: A design and evidence statement. Ask LUDIA. <https://askludia.com>

IN TEXT Parenthetical: (Stark & Rostan, 2026). Narrative: Stark and Rostan (2026). Direct quotation, with the page number of the copy you are reading: (Stark & Rostan, 2026, p. X).

CHICAGO (NOTE) Beth Stark and Jérémie Rostan, LUDIA: A Design and Evidence Statement (Ask LUDIA, 2026), <https://askludia.com>.

VERSIONING This document is dated August 2026 and is intended to be revised and superseded over time. Cite the date of the copy you read, and give the date of access if the passage you rely on may change.

CITE THE TOOL SEPARATELY LUDIA. (2026). LUDIA [AI thought partner for Universal Design for Learning]. <https://askludia.com>. In text: (LUDIA, 2026). LUDIA does not save chats, so there is no permanent link to a single conversation; note the date LUDIA was used and describe the exchange in the citing author's own words.

HOW THIS WAS DRAFTED The authors wrote this statement and are responsible for all of it. An AI assistant was used in drafting and revision. The facts and the judgements are ours.

LICENCE © 2026 Beth Stark and Jérémie Rostan. This Design and Evidence Statement is licensed CC BY-NC-SA 4.0 (Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International). It may be shared, translated, and adapted for non-commercial purposes, provided the authors are credited, any changes are indicated, and the result is licensed on the same terms. Copyright remains with the authors. We chose this licence deliberately, and changed it before publication: an earlier draft carried a NoDerivatives clause, which would have forbidden translating a document about multilingual access into any language other than the one it was written in. Nobody needs to ask us to translate this. If a translation is made, mark it as a translation and say which copy it came from, and the versioning note above explains why that matters. These are the same terms as the Needs-Based EdTech Evidence Mapping, one of the four guides read here. LUDIA itself is released separately under its own licence, available at askludia.com.

References

- Basham, J. D., & Lowrey, K. A. (2025). (Re)considering universal design for learning. *Journal of Special Education Technology*. <https://doi.org/10.1177/01626434251361443>
- Boysen, G. A. (2024). A critical analysis of the research evidence behind CAST's universal design for learning guidelines. *Policy Futures in Education*, 22(7), 1219–1238. <https://doi.org/10.1177/14782103241255428>
- Boysen, G. A., Lyles, K. D., Wylie, M. M., & Ozomma, D. F. C. (2025). A review of neurological support for universal design for learning. *Policy Futures in Education*. Advance online publication. <https://doi.org/10.1177/14782103251412473>
- CAST. (2018). UDL and the learning brain. CAST. <https://www.cast.org/resources/tips-articles/udl-learning-brain-neuroscience/>

- CAST. (2024). CAST universal design for learning guidelines version 3.0. <https://udlguidelines.cast.org>
- CAST. (n.d.). Frequently asked questions. UDL Guidelines. Retrieved 15 July 2026, from <https://udlguidelines.cast.org/more/frequently-asked-questions>
- Chile, Ministerio de Educación. (2015). Decreto N° 83/2015: Aprueba criterios y orientaciones de adecuación curricular para estudiantes con necesidades educativas especiales de educación parvularia y básica. <https://especial.mineduc.cl/wp-content/uploads/sites/31/2016/08/Decreto-83-2015.pdf>
- Colombia, Ministerio de Educación Nacional. (2017). Decreto 1421 de 2017, por el cual se reglamenta en el marco de la educación inclusiva la atención educativa a la población con discapacidad. <https://www.suin-juriscol.gov.co/viewDocument.asp?ruta=Decretos/30033428>
- EdTech Quality Collaborative. (2026). Five EdTech quality indicators. SETDA. <https://www.setda.org/priorities/edtech-quality-collaborative/> Licensed CC BY-NC-ND 4.0.
- Edyburn, D. (2024). Research or policy, which lever pushes practice faster? A universal design for learning historical footnote: A response to Guy Boysen. *Policy Futures in Education*. <https://doi.org/10.1177/14782103241260902>
- España. (2020). Ley Orgánica 3/2020, de 29 de diciembre, por la que se modifica la Ley Orgánica 2/2006, de 3 de mayo, de Educación (LOMLOE). *Boletín Oficial del Estado*, 340. <https://www.boe.es/eli/es/lo/2020/12/29/3>
- ETSI. (2021). Accessibility requirements for ICT products and services (EN 301 549 V3.2.1).
- European Commission, Directorate-General for Education, Youth, Sport and Culture. (2026). Guidelines on the ethical use of artificial intelligence and data in teaching and learning for educators (Rev. ed.). Publications Office of the European Union. <https://doi.org/10.2766/7967834> Licensed CC BY 4.0.
- European Union. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). *Official Journal of the European Union*, L 119.
- European Union. (2019). Directive (EU) 2019/882 on the accessibility requirements for products and services (European Accessibility Act). *Official Journal of the European Union*, L 151.
- European Union. (2024). Regulation (EU) 2024/1689 (Artificial Intelligence Act), Articles 5, 6, and 50, and Annex III, point 3. *Official Journal of the European Union*.
- Feine, J., Gnewuch, U., Morana, S., & Maedche, A. (2020). Gender bias in chatbot design. In A. Følstad, T. Araujo, S. Papadopoulos, E. L.-C. Law, O.-C. Granmo, E. Luger, & P. B. Brandtzaeg (Eds.), *Chatbot research and design. CONVERSATIONS 2019* (Lecture Notes in Computer Science, Vol. 11970, pp. 79–93). Springer. https://doi.org/10.1007/978-3-030-39540-7_6
- Foley Nicpon, M., Allmon, A., Sieck, B., & Stinson, R. D. (2011). Empirical investigation of twice-exceptionality: Where have we been and where are we going? *Gifted Child Quarterly*, 55(1), 3–17. <https://doi.org/10.1177/0016986210382575>
- Lindroos Cermakova, A., Havinga, B., & Clary, A. (2026). Needs-based EdTech evidence mapping. European EdTech Alliance. <https://www.edtecheurope.org/publications/needs-based-evidence-mapping-report> Licensed CC BY-NC-SA 4.0.
- Meyer, A., Rose, D. H., & Gordon, D. (2014). *Universal design for learning: Theory and practice*. CAST Professional Publishing.
- OECD. (2024). *Results from TALIS 2024: The demands of teaching*. OECD Publishing.
- Phillipson, S. N. (2008). The optimal achievement model and underachievement in Hong Kong: An application of the Rasch model. *Psychology Science Quarterly*, 50(2), 147–172.
- Poddar, S., Koley, P., Misra, J., Podder, S., Balani, N., Ganguly, N., & Ghosh, S. (2025). Brevity is the soul of sustainability: Characterizing LLM response lengths. In *Findings of the Association for Computational Linguistics: ACL 2025* (pp. 21848–21864). <https://aclanthology.org/2025.findings-acl.1125/>
- Precedence Research. (2025). Educational technology market size, share and trends. <https://www.precedenceresearch.com/press-release/educational-technology-market>
- Project Zero. (n.d.). Thinking routines toolbox. Harvard Graduate School of Education. Retrieved 15 July 2026, from <https://pz.harvard.edu/thinking-routines>

- Rostan, J., & Stark, B. (2023, September 27). Meet LUDIA, your AI-powered UDL partner. *The International Educator*. <https://www.tieonline.com/article/3575/meet-ludia-your-ai-powered-udl-partner>
- Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askill, A., Bowman, S. R., Cheng, N., Durmus, E., Hatfield-Dodds, Z., Johnston, S. R., Kravec, S., Maxwell, T., McCandlish, S., Ndousse, K., Rausch, O., Schiefer, N., Yan, D., Zhang, M., & Perez, E. (2023). Towards understanding sycophancy in language models. *arXiv*. <https://arxiv.org/abs/2310.13548>
- Siegle, D. (2018). Understanding underachievement. In S. I. Pfeiffer (Ed.), *Handbook of giftedness in children: Psychoeducational theory, research, and best practices* (2nd ed., pp. 285–297). Springer.
- Stark, B., & Rostan, J. (2026). Engaging with LUDIA through the 4 T's. Book Creator. https://read.bookcreator.com/4iqm3fqjI4RzLVdglFM2Wpl4Ei1/CxyiC_lTgSc6OSLYD-U6A (Original work published 2023)
- UNESCO. (1994). The Salamanca statement and framework for action on special needs education.
- UNESCO. (2020). *Global education monitoring report 2020: Inclusion and education: All means all*.
- UNESCO. (2021a). Recommendation on the ethics of artificial intelligence.
- UNESCO. (2021b). Windhoek+30 declaration: Information as a public good.
- UNESCO. (2023). Guidance for generative AI in education and research.
- UNESCO. (2026). *AI for environment and ecosystems toolkit for policymakers*. UNESCO Publishing.
- UNESCO & University College London. (2025). *Smarter, smaller, stronger: Resource-efficient generative AI and the future of digital transformation*. UNESCO Publishing.
- UNICEF. (n.d.). Children with disabilities. UNICEF Data. Retrieved 15 July 2026, from <https://data.unicef.org/topic/child-disability/overview/>
- United Nations. (2006). Convention on the rights of persons with disabilities, Articles 2, 9, 24, and 27.
- United Nations. (2015). Transforming our world: The 2030 agenda for sustainable development (Sustainable Development Goal 4).
- United Nations Children's Fund. (2026). *EdTech for good framework* (Version 2.0). UNICEF. Licensed CC BY 4.0.
- United Nations Committee on the Rights of Persons with Disabilities. (2016). General comment No. 4 (2016) on the right to inclusive education (CRPD/C/GC/4).
- United Nations Committee on the Rights of the Child. (2021). General comment No. 25 (2021) on children's rights in relation to the digital environment (CRC/C/GC/25).
- United States. (2008). Higher Education Opportunity Act, Pub. L. No. 110–315, § 103.
- United States. (2015). Every Student Succeeds Act, Pub. L. No. 114–95.
- W3C. (2023a). Accessible rich internet applications (WAI-ARIA) 1.2.
- W3C. (2023b). Web content accessibility guidelines (WCAG) 2.2.
- West, M., Kraut, R., & Chew, H. E. (2019). I'd blush if I could: Closing gender divides in digital skills through education. UNESCO and EQUALS Skills Coalition. <https://unesdoc.unesco.org/ark:/48223/pf00000367416>

Independent references to the earlier LUDIA

A full, searchable collection of the research, articles, videos, and podcasts referencing LUDIA, with an APA citation for each, is maintained under "Learn About LUDIA" at askludia.com. It shows that people made and published work about LUDIA. What it does not show is that LUDIA worked.