

AI assistants and neurodivergent youth

Designing responsible, inclusive and supportive AI tools for young people with diverse cognitive and intellectual abilities

Design considerations for creating responsible, inclusive and supportive AI assistants.

Part of the European Citiverses Uniting for Inclusiveness project, co-funded by the European Union

Produced by Our Normal Association as part of the European Citiverses Uniting for Inclusiveness project. The report is authored by Jenny Lindström Beijar and Anna Kolback. It presents ten themes and related design considerations for creating responsible, inclusive and supportive AI assistants, with a particular focus on neurodivergent young people and the perspectives of parents and caregivers. For use in the work of designing, procuring and regulating AI tools.

European Citiverses Uniting for Inclusiveness is an EU co-funded project exploring how Local Digital Twins and AI can be developed and used with accessibility, inclusion and citizens' needs at the centre.

We're also happy to share our findings through a presentation or workshop and discuss what responsible and inclusive AI design can look like in practice. Feel free to get in touch if you'd like to explore how these insights could be relevant for your organisation. Contact details can be found on the last page.

QUALITY OF INFORMATION

Disclaimer: Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Digital Europe Programme. Neither the European Union nor the granting authority can be held responsible for them.

STATEMENT OF ORIGINALITY

This report contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both.

In keeping with the transparency this report itself calls for, we disclose that AI tools were used throughout this project to support the literature search, and to help draft, edit, and refine sections of the text. All analysis, interpretation, and conclusions remain the authors' own, and every AI-assisted output was reviewed, verified, and revised before inclusion.

PART I — CONTEXT AND APPROACH

Conversational AI is changing what is possible. For many people, that change is subtle - a faster way to search, a more convenient way to draft an email.

For neurodivergent young people, including young people with intellectual disabilities, it can be something more fundamental: access to information they could not reach before, support that never runs out of patience, a space to ask the questions that feel too difficult to ask a person. As one parent in our research put it simply: it makes my child's world bigger.

That potential is real. It is also, at present, not fully realised, as many mainstream AI tools have not been specifically designed around the needs of this group.

The conversations happening now about AI and this group tend to focus almost entirely on risk and data bias. These are real concerns, and they deserve serious attention. But a conversation only about risk is an incomplete one. Potential and risk need to be examined together, not one after the other.

We are at a moment when the trajectory of AI development is still being shaped. The decisions being made now, about how systems are built, what values they encode, whose needs they centre - will determine whether these technologies open doors for young people who are too often left out of the design process, or simply pass them by. This report is written in the belief that those decisions matter, and that getting them right requires more than good intentions.

The gap between framework and practice

Policymakers and regulators have sought to respond rapidly to the ethical and societal challenges associated with AI. The EU AI Act explicitly identifies vulnerable groups as a priority concern. National agencies across Europe have published guidelines for responsible AI development. In Sweden, for example, DIGG and IMY have issued national guidelines for the use of generative AI in public administration (DIGG & IMY, 2025), and AI Sweden's Responsible AI Knowledge Hub offers guidelines and tools to help organisations develop and use AI responsibly (AI Sweden, 2025). Similar frameworks, guidelines, and policies exist in many other countries. The language of inclusion is present. However, regulatory and policy frameworks continue to evolve alongside rapidly developing AI technologies.

And yet, when it comes to neurodivergent young people specifically, very little of this has translated into practice. This is, by most measures, one of the most potentially vulnerable user groups in any digital context - and one of the least studied. The academic literature on

conversational AI for this population is thin. The number of products genuinely designed for their needs is smaller still.

The standard industry response to this kind of concern is age-gating: restrict access, add a disclaimer, move on. That is not enough. For many young people in this group, AI assistants are already part of daily life - used for homework, for information, for company, for processing feelings that are hard to put into words. The question is not whether they will encounter these technologies. It is whether those technologies will be supportive.

That requires something more demanding than a minimum compliance posture. It requires designers, developers, commissioners, and organisations to engage seriously with what this user group actually needs - to follow real users through real interactions, understand the daily realities of neurodiversity and intellectual disability, and build with those realities as the starting point rather than an afterthought. The dominant design logic of the technology industry - where business models and engagement metrics too often take precedence over user wellbeing - cannot be allowed to set the terms for tools serving these users. As lawsuits against Meta for addictive platform design have made clear (State of New Mexico *ex rel. Torrez v. Meta Platforms*, 2026). While these cases do not concern AI assistants or neurodivergent users specifically, they illustrate the importance of considering how engagement-driven design may affect users who can be particularly vulnerable to persuasive or manipulative design.

What we found

This report draws on two sources. The first is a systematic review of the academic literature on conversational AI and large language models for neurodivergent youth, conducted across 138 million papers (with the support of an AI research tool) and yielding 51 studies meeting our criteria. The second is a series of qualitative interviews with parents and guardians of young people aged 11 to 20 with neurodevelopmental conditions and/or intellectual disabilities - people living with these questions every day. Grounded in empathy research methodology, examining what these families say, think, do, and feel (Gray, Brown & Macanufo, 2010), the interviews add a layer of lived, everyday detail that the literature alone does not yet provide.

Taken together, the two sources point to a dual picture. The possibilities are genuine and significant: AI assistants can strengthen learning and independence, democratise access to information, provide support with life questions that some feel are hard to raise with another person, and offer a kind of tireless, patient companionship that human support systems, however well-meaning, cannot always sustain. The risks are equally real: manipulation and over-attachment, the outsourcing of thinking, the tendency to settle for an AI answer rather than taking the next real-world step, and the ever-present danger of sensitive personal data ending up in the wrong hands.

On data and privacy specifically, parents expressed something that went beyond concern; many described a quiet resignation. They knew the risks, they did not fully trust where their child's data goes, and they felt entirely unequipped to do anything meaningful about it. "It is what it is," was a recurring note. This matters, because it points directly to where responsibility must sit. Placing the burden of data protection on individuals, and particularly on the parents and guardians of vulnerable young people, is not a solution. Effective

safeguards therefore need to be built into systems by design and by default, rather than relying solely on individual users or caregivers to configure appropriate protections.

Neither side of this picture justifies ignoring the other. The potential is too significant to step back from, and the risks are too serious to wave away.

A note on perspective

The empathy research underpinning this report is conducted from a parental and caregiver perspective: the interviews are with parents and guardians, not with the young people themselves. In some cases, parents in turn raised specific topics from these conversations with their own children - but those accounts are filtered through the parent's framing and interpretation, not elicited directly from the young people. A natural next step for this line of research would be to deepen it with observational and participatory studies that engage neurodivergent youth directly, in their own words. The literature review component draws on a broader age range, including studies of young adults as well as youth - see Methodology for the full scope.

What this report is - and what it is not

This report offers a set of design considerations - grounded in evidence and lived experience - for anyone building, commissioning, or evaluating AI tools that may be used by neurodivergent young people, including young people with intellectual disabilities. It does not blueprint a specific product, nor argue for or against any particular technology.

That said, this applies to more or less all conversational AI tools and general-purpose AI assistants - not just specialised products built for this group specifically. International research suggests that developmental disabilities are common among children, with prevalence estimated at 7.5% among children under 5 rising to 13.9% among 15–19-year-olds globally (WHO & UNICEF, 2023). In Sweden specifically, Hjärnfonden (the Swedish Brain Foundation) estimates a comparable figure of approximately 10-15% of the population having a neuropsychiatric profile, or NPF (Hjärnfonden, 2016). An estimated 2% of children aged 5–14 have an intellectual disability (WHO & UNICEF, 2023). Taken together, this represents a population of many millions of users worldwide. They are not a niche audience at the margins of a product's user base – they are already a substantial share of it, whether or not that product was designed with them in mind.

And although most mainstream AI assistants have age restrictions, many children and young people use these tools in practice. They often do so without fully understanding how AI works, what its limitations are, or how to critically interpret its responses. This means that design choices around transparency, guidance, safety and support are relevant not only for neurodivergent users, but for a much broader population of young users.

The ten themes that follow address the full arc of what good design for this group requires. Many of these principles follow the logic of universal design - what is necessary for some users is better for all. A tool that is genuinely accessible to a young person with an intellectual disability is, in most respects, simply a better tool.

Each chapter presents both what AI makes possible and where the risks lie - because both are part of the honest picture. They close with a set of concrete considerations for designers and decision-makers. We hope this report can serve as a guide for both the private and public sector alike: designers and developers building AI products, public sector bodies introducing or building their own AI assistants and public AI tools, and schools and other institutions that often don't know how to take the next step.

The parents and young people whose experiences are reflected in these pages deserve tools that are worthy of their trust. We hope this report is a small step toward making that more likely.

The report's focus on neurodevelopmental conditions and intellectual disabilities follows the same logic. These are groups whose core challenges — in processing information, communicating, and navigating a complex world — sit precisely where well-designed conversational AI can make the most difference. What serves this group well is, in large part, good design for everyone.

Terminology — AI assistants and related terms

In this report, we use AI assistant as the primary term for AI-powered systems that users can interact with conversationally through text, voice or other modalities. We use AI chat to refer specifically to the conversational interaction, and AI service when referring to the broader product or service around it. Conversational AI and LLM are used primarily when discussing the underlying research and technology. Where parents or cited studies use other terms, such as chatbot, we retain their terminology where relevant.

Universal design

The UN Convention on the Rights of Persons with Disabilities (CRPD) defines universal design as the design of products, environments, programmes and services to be usable by all people, to the greatest extent possible, without the need for adaptation or specialised design (United Nations, 2006, Art. 2). The idea is closely tied to a design philosophy sometimes called designing from the edges: rather than designing for an imagined "average" user and adding accessibility afterward, the starting point is the people most likely to be excluded. What is designed to work for them tends to work better for many others too (Holmes, 2018).

Terminology and scope

Neurodiversity and intellectual disability

Neurodiversity describes the natural variation in how human brains function, process information, communicate, and experience the world. Neurodivergent is commonly used for people whose cognitive or neurological functioning differs from what is referred to as typical.

There is no single agreed boundary around what falls under the term neurodivergent. In this report, we use it broadly to include both neurodevelopmental profiles – such as autism, ADHD, dyslexia and language disorders – and intellectual disabilities. **Throughout the**

report, when we refer to neurodivergent young people, we therefore include youth with intellectual disabilities.

We may still refer to intellectual disability explicitly where it is particularly relevant to the evidence or argument, and because intellectual disability is not always understood as part of neurodivergence. People with intellectual disabilities are also particularly underrepresented in research and design.

A diverse population

The population covered here is highly diverse. Needs, abilities and experiences vary substantially between individuals, and often from day to day. Our aim is not to treat these experiences as equivalent, but to identify design considerations that recur across them.

Scope

The report focuses on young people across this broad population. The research underpinning it draws on groups with somewhat different ages and diagnostic profiles; these are described in detail in the Methodology section.

TABLE OF CONTENTS

- PART I — CONTEXT AND APPROACH
 - 1. Introduction
 - 2. Methodology
- PART II — TEN THEMES FOR RESPONSIBLE AND INCLUSIVE AI DESIGN
 - Theme 1 — Enable personalisation and adaptation
 - Theme 2 — Support user autonomy and empowerment
 - Theme 3 — Offer multimodal and accessible interaction design
 - Theme 4 — Give non-judgmental support
 - Theme 5 — Provide structural support and simplified guidance
 - Theme 6 — Motivate and offer tireless support
 - Theme 7 — Design for natural pause and reflection
 - Theme 8 — Clarify that it's a tool — like 'a robot'
 - Theme 9 — Bridge to real-world interaction and connection
 - Theme 10 — Ethical safeguards and bias mitigation
- PART III — FROM PRINCIPLES TO PRACTICE
 - 3. Co-design: Build with, not for
 - 4. Shared responsibility: Thinking in ecosystems
- PART IV — CONCLUSION
 - 5. A design choice. Not a default.
 - Afterword / Author' reflections

Methodology

This report draws on two complementary sources of evidence.

The first is a systematic literature review. A large-scale semantic search across millions of academic papers - conducted via the Elicit research platform, which draws on Semantic Scholar and OpenAlex - identified 500 studies relevant to conversational AI and large language models for neurodivergent youth. These were screened against criteria including target age range, focus on intellectual disabilities, use of conversational AI as a primary intervention, and alignment with neurodiversity-affirming approaches. The vast majority (449) were excluded at abstract screening. The remaining 51 studies were analysed for evidence on design principles, independence outcomes, user experience, implementation considerations, and barriers. The central question guiding the review was: *“What are the evidence-based design principles and implementation considerations for conversational AI and LLMs to support the independence and agency of youth aged 16–25 with intellectual disabilities?”*

Author’ note: The decision to narrow the literature review specifically to young adults with intellectual disabilities was deliberate - a wish to focus on the group least represented in the research and, in many respects, the most vulnerable, centring independence and agency during the transition to adulthood. Because the available research was so limited, the final set of sources was allowed to include studies covering broader or adjacent populations, wherever a genuine overlap with intellectual disability existed. This turned out to be valuable rather than diluting: it surfaced substantial, directly relevant evidence on neurodevelopmental conditions more broadly, and on young people with combined diagnostic profiles.

The second source is a series of 12 qualitative interviews with parents and guardians of young people aged 11 to 20 with neurodevelopmental profiles and/or intellectual disabilities. The group represented a wide range of diagnoses, including Down syndrome, autism, mild to severe intellectual disability, ADHD, dyscalculia, dyslexia, language disorder, OCD, rare syndromes, visual impairment, and hearing impairment - reflecting the real breadth of this population rather than a narrow diagnostic slice.

The interviews covered current experience with AI and digital technology, perceived opportunities and risks, questions of safety and trust, the role of school, everyday use at home, and parents’ hopes and concerns for how these tools might develop. All interviews were conducted in Swedish. Quotes included in this report have been translated into English by the authors.

The young people whose parents we interviewed had varying levels of experience with AI chatbots - ranging from daily use as a source of genuine support, to occasional experimentation, to barely any contact at all. Services mentioned included ChatGPT, Gemini, Meta AI, Snapchat’s My AI, and Perplexity, as well as voice assistants such as Siri and Google Home. Most used AI tools privately rather than in school, where its use tended to be more limited and more focused on creative tasks such as generating images or music. Several parents noted uncertainty about the full extent of their child’s use of AI tools - a finding that is itself significant.

How this report is structured

The ten themes that follow address the full arc of what responsible, inclusive design for AI assistants aimed at this user group requires - from personalisation and autonomy, through multimodal interaction and non-judgmental support, to the ethical guardrails that must underpin everything else. Each theme presents both the opportunities offered by conversational AI and the associated risks, followed by a set of concrete design considerations. This dual structure reflects the need to consider potential benefits and risks together.

The ten themes are:

- Theme 1: Enable personalisation and adaptation
- Theme 2: Support user autonomy and empowerment
- Theme 3: Offer multimodal and accessible interaction design
- Theme 4: Give non-judgmental support
- Theme 5: Provide structural support and simplified guidance
- Theme 6: Motivate and offer tireless support
- Theme 7: Design for natural pause and reflection
- Theme 8: Clarify that it's a tool — like 'a robot'
- Theme 9: Bridge to real-world interaction and connection
- Theme 10: Ethical safeguards and bias mitigation

“It needs to be both predictable and flexible. It should start from the young person’s needs – not from what technology can do.”

- Parent of a neurodivergent young person.

PART II — TEN THEMES FOR RESPONSIBLE AND INCLUSIVE AI DESIGN

Theme 1: Enable personalisation and adaptation

A consistent finding across both the literature review and the parent interviews is that personalisation is among the most important qualities an AI assistant can have when supporting neurodivergent young people. The ability to adapt to an individual's communication style, cognitive profile, sensory sensitivities, and personal preferences is not just a nice-to-have - it is often what determines whether a tool is genuinely useful at all.

AI systems tailored to the individual can reduce cognitive load, increase clarity, and provide support in ways that feel predictable and safe. The research literature points to significantly higher engagement, improved comprehension, and greater user satisfaction when systems are personalised rather than standardised. Parents interviewed echoed this strongly - many described a wish for a system that could be set up with background knowledge about their child, so that interactions would feel genuinely fitting rather than generic.

Critically, personalisation also means respecting neurodivergent ways of thinking and communicating. Hall et al. 2025 found that a GPT-4-based career chatbot consistently imposed neurotypical language, expectations, and social assumptions on neurodivergent job seekers - even when relevant user data was available. The goal must be to adapt with the user, not to steer them toward an assumed standard.

Parents also emphasised how much a single user can vary from day to day. Abilities, focus, and emotional regulation can shift significantly depending on energy, mood, and circumstance - which makes the case for dynamic, responsive adaptation all the stronger. As one parent put it:

"It needs to be both predictable and flexible. It should start from the young persons needs - not from what technology can do."

What personalisation makes possible

Knowing the person

When a system can be configured to understand a specific user - their needs, communication preferences, and daily patterns - it can become a genuinely effective support tool rather than a general-purpose chatbot. Parents spoke about wanting to enter this kind of background context once, so that every subsequent interaction would build on it:

"Being able to optionally enter background knowledge. That the AI tool knows how this person works. I'm sure that could be incredibly helpful."

Setting the pace

For many young users, one of the most meaningful qualities of AI assistant is simply that there is no rush. Unlike human conversation, which unfolds in real time with social

expectations attached, an AI assistant that adapts to the user's rhythm removes a significant source of pressure. Several parents described this as transformative: *"With the help of AI, he can control the pace. He has time to formulate himself at his own speed - he can rephrase until he gets a formulation that helps him be understood."*

Equity in the classroom

Many parents noted that truly personalised AI support could help close a real gap in school environments, where individualised guidance is often unavailable. In their view, young people who lack strong support structures at home – or whose learning needs are not well served by standard classroom methods – could benefit especially from tools that can meet them where they are.

Support that lasts

Several parents looked further ahead, imagining an AI companion that could offer continuity of support across the full arc of their child's life - including in situations where the parent is no longer present. *"As a parent I've thought a lot about the future. I imagine he could have an AI modelled on me or my wife - one that can guide him through life when we're no longer here."* This reflects not just a practical wish, but a deeper hope: that personalised AI could eventually provide the kind of sustained, knowing support that is very hard to find elsewhere.

When knowing too much becomes a risk

The cost of sharing

The deeper an AI assistant gets to know a user, the more sensitive the information involved. Many parents expressed concern about where this data goes, who has access to it, and how it might eventually be used. *"Where does the data he shares go? Who can access it? He struggles to understand what data it even collects - and he's easily influenced. If it gives him questionable advice, I don't know if he'd be able to tell it's a machine."*

"There's a reason medical record systems are locked down. Data can be sold - to insurance companies, for example. What happens to the data we share, if it ends up in the wrong hands? What if some authority suddenly decided she no longer needs extra support?"

"Probably not good at all - because of her privacy. I don't know what happens to the things that are private to her. Whose hands does it end up in, and why? What do they do with it?"

For some parents, this uncertainty translated into a clear preference: less data, kept more tightly, for a shorter time.

"I don't think it should be saved at all. If it is, it should be in a way that belongs only to him. How it's stored and limited matters - privacy matters."

Parents' feelings on this were genuinely mixed - some felt a sense of resignation, others were more guarded - but what almost all shared was a wish for meaningful control, not just a consent checkbox.

Invisible bias in the system

Without careful design, personalisation can become a surface-level adjustment that still steers users toward neurotypical communication expectations - correcting phrasing, defaulting to formal registers, or assuming social conventions that do not apply. This is a design risk that requires active and ongoing attention, not a one-time fix.

A profile that doesn't keep up

A configuration set up on a good day may not serve the same user on a difficult one. When abilities, focus, and emotional regulation shift - as they regularly do for many users in this group - a system that relies on a static profile may fail precisely when it is needed most.

"My daughter need to learn at her own pace - and that pace is not constant."

Six things to consider

For use in the work of designing, procuring and regulating AI tools.

1. **Personalisation with purpose, not engagement.**

Adaptation should serve the user's genuine needs - clarity, pacing, sensory comfort - rather than increasing usage or time-on-platform. The design goal must be centred on the user, not the product.

2. **Respect neurodivergent communication styles.**

Systems should adapt to how the user thinks and expresses themselves, without nudging them toward neurotypical norms or implicitly correcting their style.

3. **User-controlled personalisation.**

Users and caregivers should be able to adjust preferences - tone, pace, level of support, modality - and have a clear understanding of how the system adapts over time.

4. **Privacy-by-design and data minimization.**

Collect only what is genuinely necessary. Store it securely, and give users and caregivers transparent, meaningful control over what is saved, shared, or deleted - particularly for sensitive information such as health data, behavior, or emotional content.

5. **Dynamic adaptation across varying daily capacity.**

Personalisation must account for day-to-day variability within the same user. The system should be able to respond to shifts in communication, capacity, or emotional state - not rely solely on a static profile.

6. **Caregiver involvement with clear boundaries.**

Allow caregivers to contribute context and have appropriate oversight, while safeguarding the user's privacy and sense of autonomy. Transparency about what caregivers can and cannot see is essential.

"These services have given my daughter incredible independence. It feels like we've hacked the system, if you know what I mean."

- Parent of a neurodivergent young person.

Theme 2: Support user autonomy and empowerment

For many neurodivergent young people, independence is not always a given - it is something that often has to be built, piece by piece, with the right support in place. And for many in this group, the support they've been offered has not been enough - because when it relies on human resources, it quickly becomes resource-heavy. This is where AI can play a role.

A recurring theme in the interviews is that AI, at its best, can be exactly the kind of support that has long been missing: patient, always available, non-pressuring, and capable of meeting the young person where they are. The research literature reinforces this picture. Choi et al. (2024) found that autistic participants expected LLM-based conversational agents to support independent management of everyday challenges. Participants anticipated that these systems could provide a non-judgmental space that encouraged a more proactive approach to day-to-day issue (although they also expressed concerns about response quality and privacy). Jang et al. (2024) found that 82% of autistic participants preferred interacting with an LLM over a human comparison condition for workplace communication support. Participants valued the LLM for its structured responses, privacy, and the freedom to ask questions without fear of judgment or reprisal, giving them a greater sense of control when navigating workplace communication.

Koegel et al. (2025) conducted a randomised clinical trial that demonstrated statistically significant improvements in empathetic responses among autistic adolescents and adults using an LLM-based intervention. Some participants also reported modest improvements in self-confidence in targeted areas, though results on this measure were mixed. It is worth noting that this was an elective, goal-directed programme - a different context from an assistant that defaults to correcting communication style unprompted.

For this user group in particular, the point about self-confidence matters enormously. Self-confidence is not a byproduct of AI use - it can be the whole point. When a young person successfully navigates something they could not manage before, something shifts. Several parents in our interviews used the word independence almost with surprise, as if they had not quite expected to be saying it.

The examples of increased independence varied widely. One parent described how their teenager, who has ADHD and dyscalculia, could finally get explanations delivered in a way that worked for her: patiently, repeatedly, without anyone losing their temper. Another described how their son used an AI chatbot to figure out how to raise a concern with a teacher, navigating the situation entirely on his own without needing to involve his mother. These are small moments, but they add up to a significant growing sense of agency.

At the same time, this chapter carries a genuine tension at its heart. The very qualities that make AI empowering – its availability, its patience, its willingness to always provide an answer, can also make it disempowering if not thoughtfully designed. The line between increased independence and replacing it is a fine one.

How AI assistants open doors to independence

A first step into information

For many young people in this group, accessing information independently has always been difficult. Search engines require literacy, source evaluation, and the ability to filter through noise. AI tools can change this in a meaningful way. They can respond directly, in plain language, to whatever question is asked. Several parents described this as a quiet revolution in their children's daily life.

“She has never been able to Google for answers - that requires being able to read sources. With ChatGPT she can take in information in a completely different way. She can now ask about her foot condition and actually get an answer she couldn't find on her own.”

“It's great for children with intellectual disabilities because it's so accessible - she can find out things herself that she otherwise couldn't possibly find out without us. Whatever it may be. There is enormous value in that.”

Finding a voice

For young people who struggle to express themselves - whether due to language impairments, anxiety, or difficulty formulating thoughts in the moment - AI assistants offer a space to try without the social stakes of a real conversation. Several parents described their children using AI chats to rehearse or prepare: figuring out what they want to say before they say it, drafting messages to teachers, or simply practising putting words to feelings.

“He has speech, but it's quite sparse. A good vocabulary but difficulty expressing himself. Perhaps AI could generate more for him, like speak for him when he finds it hard to find the words himself.”

“He made his own checklist of notes as preparation for a counsellor visit. He came up with the idea himself using an AI chat tool. He felt satisfied and more independent. The counsellor was very pleased.”

Self-confidence through doing

Several parents pointed to an effect on their child's self-esteem. When a young person managed a task on their own - explaining something, solving a problem, handling a situation - it seemed to make a difference. The word “self-confidence” came up unprompted in a number of interviews, and parents connected it to their child's use of AI tools.

“These services have given my child incredible independence. It feels like we've hacked the system, if you know what I mean.”

“AI tools help my child interact with other people, keep track of things, and navigate the day. It makes him more independent.”

Reducing friction at home

A less visible but important benefit is what several parents described as a reduction in everyday conflict. When an AI assistant can take over the role of reminding, prompting, or explaining - tasks that often create tension between parent and teenager - the relationship benefits.

“She gets tired of us when we remind her about her morning insulin, for example. It would be great if an AI assistant did that instead. She could become more independent. Less friction, fewer conflicts at home.”

Levelling an unequal playing field

Autonomy enabled by AI assistants is also, quietly, an equity issue. For young people who lack strong support structures at home, or who come from socioeconomically disadvantaged backgrounds, AI can be a meaningful equaliser - offering the kind of patient, structured support that better-resourced peers simply take for granted, from a resource that asks nothing social in return. As one parent put it:

“AI chatbots can support a child with neurodevelopmental conditions - for example one who struggles with planning and follow-through. You can get a completely different kind of structural support - the kind another child might get from a capable parent at home. For children who are socioeconomically disadvantaged and/or whose parents’ might have neuropsychological diagnoses themselves, that makes a real difference.”

When support slides into dependency

The answer that goes nowhere

One of the more nuanced risks that parents raised is not that AI chats give wrong answers, but that they give complete answers - and that can be enough to stop a young person from taking the next step. For someone with ADHD, the cognitive relief of having received a response can feel like the task is done, even when it isn't.

“A potential risk is settling for an answer without taking further action. For individuals with ADHD, initiating tasks can be difficult, so they might avoid following up such as booking an appointment because they feel they’ve ‘already gotten an answer.’ Since there’s often a tendency to choose the easiest path, important next steps may be skipped.”

Thinking that gets outsourced

Several parents expressed a worry that mirrors a broader debate now unfolding in both research and policy: the concern that their child will stop developing their own thinking if AI assistants are always there to do it for them.

“You get a bit of brain rot, there’s something to that. She should learn herself. I want her to think a little on her own sometimes, not just let the AI tool think for her. And this applies to me too, of course!”

A tool that doesn't know what it doesn't know

AI assistants today are consistently confident in tone, regardless of whether their answer is correct, partial, or context-inappropriate. For young users who already struggle with critical evaluation, and who may have no frame of reference to detect when something is wrong, this is a real hazard. The problem is compounded when no sources are shown and the response appears authoritative. Understanding that AI tools can be wrong is difficult for many adults; for this user group, it can be genuinely impossible in areas where they lack background knowledge to cross-check.

The fine line

Perhaps the most honest summary came from a parent who has seen both sides:

"The line is very fine - you become enslaved to it if you don't understand what it does. That's a real risk. You become less independent if you're dependent on it to get through the day. That said, I think it's a balance worth exploring and testing. Something that gradually makes individuals more independent would be good."

Five things to consider:

For use in the work of designing, procuring and regulating AI tools.

1. Support action, not just answers

AI should be designed to help users recognise when they have enough to act - and feel confident doing so - rather than defaulting to offering another round of interaction. The measure of success is the user taking a real step, not continuing the conversation.

2. Guide thinking rather than replace it

Rather than providing immediate full solutions, the system should offer hints, guiding questions, and step-by-step prompts that keep the user in the driving seat. The goal is to strengthen independent thinking over time, not substitute for it.

3. Make uncertainty visible

AI tools should clearly and accessibly signal when an answer may be incomplete, context-dependent, or uncertain - and do so in language the user can understand. Confident-sounding responses to questions in sensitive areas (health, safety, relationships) carry particular risk for this user group.

4. Build in a bridge to human support

Design interactions that naturally and warmly guide users toward involving other people when appropriate; a teacher, a caregiver, a professional. The transition should feel supportive rather than like a barrier or rejection. (See also Theme 9: Bridge to real-world interaction and connection, which explores what healthy transitions from AI to human support look like in practice.)

5. Design for gradual, adjustable support

The level of support offered should be adjustable over time; decreasing as the user builds confidence and capability, or increasing when they are having a harder day. A good support is one that eventually lets you stand on your own.

"She started by writing, but now she speaks into it. It answers back with both voice and text. That way she manages the chat very independently."

- Parent of a neurodivergent young person.

Theme 3: Offer multimodal and accessible interaction design

How a person communicates is not only a preference – for many it is often a necessity. For many neurodivergent young people, the channel through which information arrives, and the form it takes, can be the difference between understanding and exclusion. A text-heavy interface that works well for a neurotypical user may be entirely inaccessible to someone who cannot read fluently, processes information better through images, or finds typing painfully slow. This is not a niche consideration. It is the starting point for any design that claims to serve a broad population.

Multimodal interaction; combining voice, text, images, pictograms, symbols, and audio, is now broadly recognised as essential for inclusive AI design. Xu et al.'s (2025) scoping review of 117 papers found that multimodal interaction, adaptive feedback, and embodied interfaces can enhance engagement and usability for neurodiverse users. Yet the review also found that most studies fail to adequately address sensory and cognitive heterogeneity. Taken together, these findings suggest that the gap between what is technically possible and what is actually built remains wide.

What makes this chapter distinct from Theme 1 (personalisation) is its focus on the interaction layer itself - not just what the system knows about the user, but how it communicates, in what form, and whether that form can shift. A system can be perfectly personalised in its content and still fail completely if its interface requires the user to type fluently, parse long sentences, or interpret abstract icons.

Parents in our interviews gave a vivid picture of what this looks like in practice. Children and young people found their own routes into AI tools; discovering text-to-speech, switching to voice input, or finding that having answers read aloud made a conversation possible when reading did not. These were not guided choices. They were workarounds found by users themselves, which says something important about both the potential and the current state of design.

"She can't read, so she always has the answers read aloud to her. She found that herself - I hadn't actively shown her."

"She started by writing, but now she speaks into it. It answers back with both voice and text. That way she manages the chat very independently."

A Vinnova-funded project collaboration between RISE and the Swedish municipalities of Skellefteå, Gothenburg, and Jönköping, developing an AI-based communication support prototype in close collaboration with people with intellectual disabilities (Sjölinder et al., 2025) provides valuable practice-based evidence on this topic. Workshop participants were clear and specific about what they needed: voice input, text read aloud slowly and clearly, images alongside text, short direct answers, and the ability to adjust the level of simplification. The

project also found that offering too many settings was counterproductive; the right balance between flexibility and simplicity is itself a design challenge.

The channels that open doors

Images and pictograms as a first language

For users with severe communication needs or limited literacy, images are not a supplement to text - they are the primary medium. Saffar et al. (2024, 2025) conducted two studies specifically with people with intellectual disabilities on image-based self-expression, finding that LLMs show encouraging capabilities for detecting aspects of communicative intent from image selections, and that AI-guided image recommendation can support non-verbal communication. Given the limited research applying these approaches specifically to people with intellectual disabilities, these findings are particularly relevant. The Vinnova-funded AI-based communication support project (2024–2025) tested pictogram generation as a core output modality, and workshop participants asked for visual supports as a standard option, not an add-on. Parents in our interviews also requested this:

“I’d like more qualified image support for my son - more like a unified language that chatbots can generate.”

“Image support for language is a good feature. A chatbot that could interpret my child would be great; one that understands signs, how my child speaks, and can mirror through images. It could hold a longer thread in a conversation, help avoid changing topics constantly and both follow along and lead.”

Speaking instead of typing

For many users in this group, typing is a significant barrier - whether due to literacy difficulties, motor challenges, or cognitive effort of formulating text. Voice input can support removing that barrier. Workshop participants in the Vinnova-funded AI-based communication support project (2024–2025) consistently requested the ability to speak their questions. Several parents in our interviews also described their children as having discovered this route independently. The potential is clear, but important limitations remain. Current systems still struggle with atypical speech patterns. Pucci et al. (2024) found that people with speech impairments, specifically dysarthria, reported that current conversational AI systems do not adequately meet their needs - a significant gap given the potential importance of voice input for this population.

Cha et al.’s (2021) study of adolescents with autism using voice-based conversational agents found that these systems could be highly engaging and empowering, but also highlighted the need to accommodate users’ diverse needs, personalities, and expectations.

“My son has a visual impairment, and in the afternoon he has more difficulty seeing, so then he prefers to speak more to his devices.”

“My children have spoken to their devices, or let the AI speak. They’re not fast at reading and writing. For them it’s quicker to speak and listen to answers.”

Text that speaks and simplifies

For users who can engage with text but need it transformed – for example, read aloud, simplified, or broken into shorter pieces, AI offers capabilities that previously required specialist tools. The Vinnova-funded AI-based communication support project (2024-2025) found that text-to-speech was a core need, and that the read-aloud function needed to be slow and clearly paced to be genuinely useful. Jerry et al. (2026) developed a model-driven architecture that dynamically adapts interface content and presentation, producing plain-language text, pictograms, and high-contrast layouts aligned with established accessibility standards and guidance, including WCAG 2.2 and W3C COGA. The Vinnova-funded AI-based communication support project (2024-2025) similarly built multiple levels of text transformation; from simplified language to bullet points to image support, as switchable output modes.

“It would be fantastic help in school too! For someone who is dyslexic, getting things read aloud, summarising. It feels like it’s been kept away - as if it doesn’t exist. But it does, and it will only become more present in different environments.”

Conversation as navigation

Several parents and workshop participants in the Vinnova-funded AI-based communication support project independently identified something important: many users in this group find keyword-based search deeply frustrating, but can manage a conversation. Being able to ask a question in plain words, in their own phrasing, and receive a response that matches their tone is a fundamentally different experience from trying to construct a search query. The Vinnova-funded AI-based communication support project found that workshop participants strongly preferred a dialogue-based interaction style and wanted AI to respond in the same register as the question. A friendly, consistent tone was not a nice-to-have; it was what kept users engaged and willing to try again.

“Images, film, repetition - that suits my daughter.”

When the channel breaks down

Voice that doesn’t understand

The promise of voice interaction depends entirely on whether the system can understand the user. For people with speech impairments, non-standard pronunciation, or atypical speech patterns - all common in this population - current recognition systems regularly fail. Pucci et al. (2024) found this was a fundamental unmet need among people with dysarthria. Parents in our interviews described the same problem directly.

“We had Alexa before but it had difficulty understanding my daughter because she has a speech impediment. It wasn’t good at that then. I feel the services understand better now, but these LLMs need to be trained on different ways of speaking.”

“Sometimes the chat doesn’t quite hear what she says - and then it answers something completely different. In those cases, if it’s a topic she knows well, she can correct the AI. But if it had answered in medical language about her foot condition and said something wrong - she wouldn’t have been able to understand that.”

This last observation carries a double risk because of the invisibility of the failure mode: misrecognition leads to a wrong answer, and the user may have no way to detect the error.

Sign language; the missing modality

Despite being a primary language for many users, sign language remains poorly supported by current AI systems. Cheng et al.'s (2024) study of d/Deaf and hard-of-hearing learners identified the lack of sign-language support in current LLMs as an explicit accessibility gap. Catania et al.'s (2023) systematic review similarly highlighted the importance of multimodal interaction, including gestures and signs, in conversational agents for people with neurodevelopmental disabilities. Recent developments suggest this may be starting to change, with major AI labs beginning to invest in dedicated sign-language translation models - though such capabilities remain far from standard in general-purpose conversational AI systems. This is not a fringe need - it represents a fundamental accessibility gap in the current generation of tools.

Customisation that overwhelms

Offering multimodal options is not enough if accessing them requires navigating a complex settings interface. The Vinnova-funded AI-based communication support project (2024–2025) encountered this directly: initial prototypes with extensive customisation options proved difficult for users, and the team eventually settled on three levels; a simple persona preset, in-chat adjustment buttons, and an advanced sidebar for those who wanted it. The insight is generalisable: flexibility must be offered in a way that is itself accessible. A user who needs image support but cannot find the setting has not been helped. Parents in our interviews highlighted this as well:

“You have to simplify a lot. The interface must be very easy to use. Easy to get into. Not adding pop-ups everywhere.”

The equity gap

Multimodal features can themselves create barriers when they depend on particular devices, connectivity, or levels of digital literacy. Xu et al.'s (2025) scoping review similarly highlights persistent accessibility barriers and the failure of many existing approaches to accommodate sensory and cognitive heterogeneity. AI tools should therefore be designed for flexible access across devices and contexts, and for use both independently and with support from others around the user.

Four things to consider:

For use in the work of designing, procuring and regulating AI tools.

1. Design for the full range of communication

Systems should be built from the start to support voice input, text-to-speech, image and pictogram output, plain-language transformation, and dialogue-based interaction; not as add-ons, but as core modalities.

2. Make modality switching simple and visible

Users should be able to move between interaction modes easily and intuitively, without navigating complex settings. Give users a simple default and let them go deeper only if they want to.

3. **Train for atypical speech and communication patterns**

Voice recognition systems must be developed and tested with users who have speech impairments, non-standard pronunciation, or atypical pacing. A system that only works for fluent, neurotypical speech is not accessible. Sign language support, currently absent from virtually all AI systems, should be treated as a priority rather than a future consideration.

4. **Follow accessibility standards - and push beyond them**

WCAG 2.2 and W3C COGA guidance provide an important baseline for cognitive and sensory accessibility in digital interfaces. However, these standards were developed primarily for web content and their application to conversational AI interfaces remains incomplete. Designers should treat them as a floor, not a ceiling, and actively seek guidance from co-design processes with neurodivergent users to fill the gaps.

“When you have a neurodevelopmental condition, you often feel different – and that’s something you may need support with. Especially things that are difficult to talk about with others.”

- Parent of a neurodivergent young person.

Theme 4: Give non-judgmental support

There is a particular kind of exhaustion that comes with feeling constantly evaluated. For many neurodivergent young people, ordinary interactions can carry that weight; asking a question in class, raising a concern with a teacher, trying to explain a feeling to a parent who is already stretched thin, or simply wanting to know something about their own body without the conversation becoming a whole thing. The social cost of asking is often high enough that many questions simply go unasked.

This is one of the quieter but more significant things that AI assistants change. A system that responds without frustration, without judgment, without the subtle signals of impatience or surprise - and that is available at any hour, on any day - removes that social cost. Questions can be asked, feelings can be named and situations can be rehearsed before they have to be faced in real life.

Jang et al. (2024) found that 82% of autistic participants preferred interacting with an LLM over a human providing the same type of workplace communication support, valuing in part the freedom to ask questions and explore situations without fear of judgment or social consequences. Choi et al. (2024) found that autistic individuals saw LLM-based conversational agents as offering a non-judgmental space that could encourage them to address day-to-day challenges proactively and independently. These findings resonate with something parents interviewed for this study described just as clearly: their children used AI chats to talk about things they might not bring to another person, giving them a space to work through thoughts and concerns that might otherwise remain unspoken.

At the same time, several parents described a more cautious concern - that this same lack of judgment could tip into something else, an AI chat becoming “too much of a friend” - a tension this report returns to in Theme 9 (bridging to real-world interaction).

What this chapter is specifically about - and what distinguishes it from Theme 2 (autonomy) and Theme 9 (bridging to real-world interaction) - is the quality of the interaction itself. Not what the user can accomplish, but how it feels to ask. Non-judgmental support is a design principle, not just a byproduct of AI’s nature. It requires active choices about tone, consistency, response to emotional content, and the limits of what the system should and should not say.

A space to ask the questions that feel too hard

The questions that don’t get asked anywhere else

For many young people in this group, there is a long list of things they want to know but cannot easily ask; about their diagnosis, their body, their feelings, their fears. These are not unusual questions for adolescents, but the usual routes to answers, such as parents,

teachers, search engines all carry social or practical barriers. AI tools can lower those barriers. Parents described their children using chatbots to explore health concerns, ask about puberty, process existential worries, and try to understand what their diagnosis means.

“She wrote about her diabetes and her fear of dying once - that was valuable because she hadn’t dared to bring that to us as parents. AI chats can be a support there.”

“He has a lot of existential thoughts - like how people work and how babies are made. I think he uses it to ask things he finds awkward to ask us. For example, how to make friends. He’s quite lonely and looks for guidance on life questions.”

“Some questions about sex, for instance are very hard, even though many people have the same ones. It can feel uncomfortable to ask in real life – and then it gets easier in another forum.”

Emotional rehearsal and self-expression

Several parents described AI assistants as a space where their child could practice putting words to feelings which is something that for many in this group is a timely and difficult process. Being able to try out how to express something, without the pressure of another person’s reaction, can make the real conversation easier when it comes. Choi et al. (2024) similarly found that autistic participants saw conversational AI as a way to explore communication strategies and rehearse how they might navigate social situations. Parents interviewed for this study raised similar aspects:

“We’re working a lot with him on putting words to his feelings.”

“Right now both he and we struggle to understand all the feelings. Which are ADHD, which are hormones and puberty? A teenager wants to ‘switch off’ parents. But something who can listen, be a small support, a bridge - that’s still needed when you’re handling a lot of simultaneous feelings.”

Consistent, tireless, never sighing

One theme that ran through the parent interviews with quiet force was the simple reality of parental fatigue. Parents love their children, but they also get tired. A child who asks the same question forty times, or who needs the same explanation repeated in eight different ways, is not doing anything wrong. But human patience has limits, and AI tools don’t. Several parents described this honestly and without guilt: the fact that AI tools never lose patience, never sounds irritated and never makes the child feel like a burden was experienced as something very appreciated.

“It’s incredibly easier to talk to AI chatbots for my son. He struggles to formulate himself. With AI chats he can control the pace and rephrase until he gets a formulation that makes him understood.”

When non-judgmental becomes uncritical

A system that always agrees

The same quality that makes AI chatbots feel safe, like its consistent warmth and its willingness to validate, is also its most significant risk in this context. AI systems are known to be prone to telling users what they want to hear. For a young person who is emotionally vulnerable, or who is processing something difficult, a response that is endlessly affirming may feel supportive while actually doing harm. Parents named this directly.

“When you have neurodevelopmental conditions, you often feel different and need support with that. Things that are difficult to talk to others about. But chatbots often say what you want to hear - that can definitely be a risk for my daughter.”

Answers mistaken for truth

When a system responds with warmth and apparent confidence, and when no sources are shown, it is very easy to believe that what it says is correct. For this user group, critical evaluation of AI responses is particularly difficult - and in emotionally loaded areas like health, relationships, or self-understanding, a wrong or incomplete answer can have real consequences. The risk is not just misinformation in the abstract; it is a young person making a decision about their health, their relationships, or their self-image based on an answer they had no way to question.

“He would not be able to do source criticism. That he wouldn’t understand.”

Settling for the conversation instead of taking the next step

Non-judgmental support can also become a substitute for the human contact that is actually needed. A young person who processes a difficult experience with an AI chat tool, and feels heard, may not then bring it to a parent, a counsellor, or a doctor - not because they don't trust those people, but because the need has been partially met. This is a design tension that requires active attention. (See also Theme 9: Bridge to real-world interaction and connection.)

Sensitive areas require clear limits

Not all questions that feel safe to ask are safe for AI tools to answer without guardrails. When conversations move into territory involving self-harm, medication, extreme emotional distress, or dangerous situations, the system needs to respond with a clear, kind boundary - and remind the user, simply and directly, that it is a tool, not a person, and that some things are better brought to a human who has real life expertise and can truly help. This is not about just shutting the conversation down, but about protecting the user by being honest about what AI chats are and aren't.

Four things to consider:

For use in the work of designing, procuring and regulating AI tools.

1. Tone as a design decision

The non-judgmental quality of an interaction has to be actively designed: consistent patience, the absence of frustration, and a stable response to unusual or sensitive questions do not happen by default. These qualities should be built in, tested with this user group, and maintained across the full range of interactions. At the same time, the system should always be transparent that it is a tool - not a friend, therapist, or trusted adult.

2. Honest, not just affirming

A system that only validates is not genuinely supportive. AI assistants should be designed to respond with care and honesty, including when the honest response is "I'm not sure" or "this is something worth talking to a person about." For emotionally sensitive topics especially, gentle honesty serves users better than endless agreement.

3. Clear, kind limits in sensitive areas

Systems must be designed to recognise when a conversation has entered territory that requires human support, and to communicate that clearly and warmly - reminding the user in simple language that AI tools have limits, and that a real person is the right next step. (See also Theme 9.)

4. Design with awareness, not ambition

Young people in this group will use AI tools to process feelings, ask difficult questions, and seek comfort - whether or not that was the designer's intention. The design should acknowledge this reality thoughtfully: maintaining clarity about what the system is, keeping responses honest and grounded, and actively supporting the distinction between AI system interaction and human connection - particularly when conversations move into sensitive territory.

“The tools that exist today are so clearly made by someone without neurodevelopmental conditions. Long checklists and lots of text – that doesn’t work for us.”

- Parent of a neurodivergent young person.

Theme 5: Provide structural support and simplified guidance

For many neurodivergent young people, the challenge is not understanding something, but rather about getting started, keeping track of the steps, and not losing the thread when the whole feels overwhelming. Executive function difficulties, working memory challenges, and low tolerance for complexity are common in this group, and they make the gap between knowing what needs to be done and actually doing it feel very wide.

This is a practical chapter about a practical problem. AI assistants can help break that gap down by turning an overwhelming whole into manageable steps, offering a structure that adjusts to the user’s pace, and providing the kind of consistent, patient guidance that is often unavailable in school or at home. Several parents described experiencing this directly themselves, as adults with their own diagnoses, before seeing the same benefit for their children.

What structure makes possible

Breaking the overwhelming into the manageable

For a young person facing a school assignment, a form to fill in, or a new routine to learn, the whole can feel paralysing before a single step is taken. Tang et al. (2025) found that transforming dense text into interactive mind maps showed promise in improving motivation, concentration, and task planning for students with ADHD. AI tools can similarly break down a larger task into a sequence of concrete, manageable steps, sized to what the user can handle right now. Several parents described this as transformative, both for their children and for themselves.

“Study planning is incredibly difficult for him. Getting an overwhelming whole broken down into small steps, and in a way that resonates with him, would be great. It doesn’t need to be a physical person. That’s what’s missing in school today.”

“For my son the instructions of an assignment must be short, concise, and adapted - but still giving all the information. Otherwise it just goes straight through.”

“I’ve taught Claude how I function. Claude asks me: ‘Should you really be doing that now - shouldn’t you finish that first?’ That’s really good for me, and I can see that something similar would benefit my daughter as well.”

Guidance that fits the person, not the tool

Several parents noted that existing tools, including those provided by habilitation services, were clearly designed by people without lived experience of neurodevelopmental conditions. Long checklists, dense text, and rigid structures miss the point entirely. AI assistants offer

something different: the ability to adapt the level of structure, the length of response, and the degree of simplification to the individual - and to change that in real time.

“The tools that exist today are so clearly made by someone without neurodevelopmental conditions. Long checklists and lots of text - that doesn’t work for us.”

“We have a digital screen from habilitation that structures his days - he uses it. But imagine if that screen could also talk to him. The one we have today is very complicated to use. And ugly.”

Where structure can fall short

Oversimplification that stops too soon

There is a risk in making everything maximally simple all the time. Some users need the option to grow, to be challenged a little more as they build confidence and to move from step-by-step guidance toward more independent functioning. A system that only ever simplifies may unintentionally limit what a user believes they can do.

Structure without follow-through

Getting started is one thing; completing and following through is another. AI tools can help a user plan beautifully and still leave them without the support to actually execute. There is also a real risk of a gap between a user’s perceived competence and their actual level - something parents described directly:

“Sometimes he thinks he knows everything. And then he doesn’t understand why he gets a specific grade in school. His sense of his own knowledge level and his actual knowledge level - that’s hazy.”

Designed by people who don’t need it

As parents observed, many existing tools fail this group because they were built without genuine co-design. The same risk applies to AI tools. Structural support that works for a neurotypical user - with its assumptions about attention span, reading ability, and tolerance for complexity - may be actively unhelpful for the people who need it most.

Three things to consider:

For use in the work of designing, procuring and regulating AI tools.

1. Make the level of detail adjustable

Users should be able to move between small, concrete steps and a more complete overview, depending on their energy, capacity, and the task at hand. When needed, the system should make it easy to focus on one step at a time without losing access to the bigger picture. This is about the amount of information shown at once, not about how much the system should encourage independence (which is discussed in theme 2).

2. Let the user choose between educational and simplified mode

Separate from how much is shown is whether the system tries to build skill over time. An "educational mode" can nudge the user a little further with each interaction - prompting them to attempt the next step themselves rather than handing it over. A "simplified mode" does the opposite: it gives the answer, plainly, without a lesson attached. Which mode is right depends on the day, the task, and the user's energy - and the choice should belong to them/be possible to select in settings.

3. **Design for follow-through, not just planning**

Structural support should extend beyond breaking things down to helping users stay with a task; confirming steps get done, and gently keeping the thread. The goal is completing, not just starting.

“My child needs to be able to stay with something longer and go deeper – not just switch to the next thing all the time and lose interest.”

- Parent of a neurodivergent young person.

Theme 6: Motivate and offer tireless support

Human patience and resources have a natural limit. Parents get tired, teachers have thirty other children, support staff have schedules... For young people who need to ask the same question many times, who learn through repetition, who need encouragement to get started and gentle persistence to keep going - that limit matters.

“When he gets over the threshold and wants information, I think an AI chat could be a golden tool - especially if he can use his voice, then he avoids unnecessary detours. He’s extremely impatient. Gives up in two seconds if there’s something he doesn’t get.”

“My child needs to be able to stay with something longer and go deeper – not just switch to the next thing all the time and lose interest.”

What tireless availability makes possible

Repetition without friction

Many young people in the target group learn through repetition - asking the same thing in slightly different ways, returning to the same topic, needing the same explanation reframed until it lands. AI assistants can respond without the subtle cues that might make a child feel like they’re asking too much.

Joy, play, and curiosity

Several parents described their children using AI tools in ways that were neither task-focused nor emotionally heavy - simply just playful. The examples ranged between using AI chats as an audience for imaginative games, as a source of random facts to something fun to interact with. This dimension of engagement is easy to overlook in a design context, but it matters: a tool that is also enjoyable to use will be used, and for young people who may have limited social outlets, a source of low-stakes, joyful interaction has real value.

“Sometimes she plays that she is the artist and the stuffed animals and the AI chatbot are the audience. She uses it in play - but doesn’t think of it as a human. She’s never asked ‘shall we meet up?’ It’s more that she asks questions and involves it in her play.”

“My kids often use AI chats just as a fun thing. Grandma has a Google home device - they think it’s hilarious and often ask things like ‘can you fart?’”

Filling a gap when human support runs thin

Several parents spoke honestly about moments when they had run out of energy - and described AI chats as something that could carry the load in those moments without

judgment on either side. They didn't describe it as a replacement for human connection, but as a meaningful complement:

“Our daughter gets tired of us, and we’ve got a little tired too. We can’t always hook onto everything. Then it’s good to have something that always has the energy to answer.”

When availability becomes a problem

The same quality that makes many existing AI assistants tireless – their willingness to continue – can become a risk for users who struggle with self-regulation and transitions. A conversation that never has to end, a game with no natural stopping point, an interaction that is always just one more prompt away: for some users, this is not a feature but a trap.

This risk is closely connected to the design principles in Theme 7 (Design for natural pause and reflection), which addresses it in full. The key point here is that motivating design and addictive design are not the same thing, and the line between them matters enormously for this user group.

The interviews highlighted a specific concern worth noting: children will test limits, and a system that has no clear boundaries around things like content, around duration, around what kinds of interactions it will sustain - is a system that will eventually fail the user it was meant to support.

Four things to consider:

For use in the work of designing, procuring and regulating AI tools.

1. Design for patience, not performance

The tireless quality of AI assistants should be expressed through consistent availability and willingness to repeat - not through escalating engagement techniques. Motivating design means staying with the user; it does not mean keeping them hooked.

2. Take play seriously

Playful, low-stakes interaction is a use case to take into consideration as well. Systems should be able to engage with curiosity, humour, and imagination without treating every interaction as a task to be completed. At the same time, even playful interactions should maintain the same clear, honest tone about what the system is.

3. Build in natural endings

Systems should be designed with clear, gentle ways for interactions to conclude - particularly for users who find transitions difficult. An interaction that is easy to start must also be easy to stop. (See Theme 7 for the fuller design principles around pause and transition.)

4. Give caregivers meaningful controls

Parents and support staff should be able to set boundaries around duration, topics, and intensity of use - not as a restriction on the young person's autonomy, but as a

responsible part of the overall design. These controls should be simple to use and clearly explained. How they are configured will depend on the individual user's needs and capacities. (For guidance on personalised setup, see Theme 1. For the broader framework of ethical guardrails that should govern what any system permits, regardless of individual settings, see Theme 10.)

“The Pippi game on my son’s ipad just keeps going and going. There has to be a stop. Now we set a timer and then the game ends – that’s to help him learn.”

- Parent of a neurodivergent young person.

Theme 7: Design for natural pause and reflection

Most digital products are designed to keep people engaged. Notifications, streaks, rewards for returning - the mechanics of attention capture are well understood and widely used. For most users, this is an annoyance. For young people with neurodevelopmental conditions or intellectual disabilities, who may already struggle with self-regulation, transitions, and disengagement, it can be genuinely harmful.

This chapter is about a design choice that runs counter to the dominant logic of the technology industry: building in endings. Natural stopping points, gentle prompts to pause, interactions that conclude rather than perpetuate. This chapter connects directly to concerns raised throughout this report about dependency, loops, and the risk of AI assistant interaction becoming something that cannot be switched off.

“The Pippi game on my sons’ lpad just keeps going and going. There has to be a stop. Now we set a timer and then the game ends - that’s to help him learn.”

“I think it’s good that she gets her dose of talking. She thinks the chat is fun, but it’s also good that it comes to a natural stop, so that we don’t have to ‘take it away’. When it ends with something like ‘It’ll be fine, I hope you have a lovely day’ - then it’s been like a conversation and a little company for a while.”

Why endings matter for this group

Transitions are hard by nature

For many young people with neurodevelopmental conditions, switching from one activity to another is genuinely difficult - not a matter of willpower but of neurological wiring. A tool that can go on forever removes a friction that served a purpose. Ending the interaction on its own offers the user a natural transition, rather than requiring them to initiate one.

“Companies shouldn’t exploit those children. Not just give them more, more, more.”

The dopamine dynamic

Several parents drew an explicit parallel with social media and gaming. Systems designed to maximize engagement use reward mechanics that are particularly potent for users with ADHD and related conditions. Some AI assistants are not yet designed this way, and parents generally felt the current risk was lower than with games or social media. But the direction of travel in the technology industry is toward more engagement, not less, and youth in our target group is particularly vulnerable to that dynamic. The concern about AI tools moving toward gamified, dopamine-driven design like daily streaks, reward animations, endless follow-up questions, was raised explicitly by many of the parents we interviewed:

“My worry is that there’s a risk if they go in the social media direction - that they steer AI chats toward giving people dopamine kicks and that algorithms make them just keep watching more. That would not be good.”

Reflection as a design goal

Beyond stopping, there is the question of what happens in the pause. Several parents raised the idea that AI tools should actively encourage users to think for themselves - not always providing the answer immediately, but occasionally asking “what do you think?”, “what have you tried?”, or simply giving the user a moment before continuing. This is a small design choice with potentially significant consequences for learning and development.

“You get a bit of brain rot - there’s something to that. She should learn herself. I want her to think a little on her own sometimes, not just let the AI LLM think for her every time.”

The risk of getting this wrong

The risk here is essentially one: that AI assistant interaction, designed without attention to natural endings and reflection, slides toward the same engagement patterns that make social media harmful for this group. It might not even be because of malicious intent, but through what has become the ordinary logic of a digital product that is “optimised for use”.

“With neurodevelopmental conditions, the dopamine from phones is hard to beat. It can become harder to stop because of that.”

Parents were clear that screen time is not inherently bad - the question is rather what kind of screen usage, and whether the user retains meaningful control over when it ends.

“Screen time is not just good or bad - it’s somewhere in the middle.”

Four things to consider:

For use in the work of designing, procuring and regulating AI tools.

1. Build in natural endings

Every interaction should have a clear and gentle way to conclude - not just a session timeout, but a designed closing moment that feels like a natural end rather than an abrupt cutoff. A simple heads-up that a break is coming soon, followed by a closing phrase, helps users who struggle with transitions prepare rather than be caught off guard.

2. Do not design for over-engagement or addiction

AI assistants should be explicitly designed against the mechanics of attention capture: default mode should not include things like streaks, reward animations for returning or algorithmically optimised follow-up questions designed to extend the session. The measure of success should be whether the user got what they needed, not how long they stayed.

3. Encourage reflection, not just response

Systems should occasionally nudge users toward their own thinking rather than always providing an immediate complete response - not necessarily as a question that demands an answer, but as a gentle prompt: "you might want to reflect on this yourself and decide whether you agree." This supports independent thinking and slows the interaction down in a way that benefits users who process at their own pace.

4. **Support caregivers in setting boundaries**

Parents and support staff should have simple, effective tools to set limits on duration and session frequency - and those tools should actually work. Several parents described feeling outmaneuvered by their child's digital literacy. Parental controls that can be bypassed are not controls. (See also Theme 1 for personalisation of these settings and Theme 10 for the ethical framework that should govern what systems are permitted to do regardless of individual configuration.)

"She understands that it's a 'robot'. Though I'm not sure whether she interprets the answer as AI generated or just as 'true'."

- Parent of a neurodivergent young person.

Theme 8: Clarify that it's a tool – like 'a robot'

Most young people, according to the parents we interviewed, knew on some level that they were talking to a machine. Parents interviewed for this study described their children calling it "the robot," "the computer," "the chat," and said they had often had deliberate conversations about it – some after seeing relevant clips on kids' TV explaining what AI is. And yet, the same parents described their teenagers taking everything the system said as true, forming emotional attachments, feeling comforted by its responses, and in some cases genuinely grieving when an interaction ended. Understanding that something is a robot and experiencing it as a robot are two very different things – and for this user group, the gap between those two can be even wider.

The question is not simply whether AI assistants should declare themselves non-human – it should, clearly and always. The question is how, in what form, and with what design consequences. Our interviews suggest that, for some neurodivergent young people, the standard solution, like a disclaimer in the onboarding, a line in the terms of service, is essentially not enough.

Research with autistic users points to an important distinction: knowing that an AI is not human does not necessarily mean knowing when not to trust it. Choi et al. (2024) found that autistic participants could find it difficult to critically evaluate responses from conversational AI, including inaccurate information. Haroon and Dogar (2024) identified a related concern in their study of TwIPS, where participants worried about becoming reliant on the system's judgements even when those judgements might be wrong. The authors argue for making uncertainty more transparent to users. For this group, then, clarity needs to extend beyond what the AI is to an understanding that what it says may not always be right.

What clarity makes possible

A functional relationship with a useful tool

When young people better understand what artificial intelligence is, they can use AI tools more effectively. They can take answers as starting points rather than verdicts. They can disagree. They can bring what they learned to a real person for confirmation. Several parents described this functional mode approvingly – their child used AI tools as a resource, not an authority:

"She understands that it's a 'robot'. Though I'm not sure whether she interprets the answer as AI generated or just as 'true'."

"I think it's better that it doesn't look so incredibly fun. A pink unicorn would have added so much for children – now the AI chat we use in our home is more like a 'robot' / a chat that gives good answers. Then she uses it more functionally."

"In the beginning my daughter thought it (the AI chat) was a 'him'. We tried to explain, but how do you say that it's 'a robot' on the other end? She immerses herself in different worlds – she probably wants them to be a real person."

Design that keeps the distinction visible

Parents in our interviews offered some specific design observations – several pointed out that visual and interaction design choices carry enormous weight in how children experience “AI identity” - far more than any verbal or written declaration. A chat interface that looks and behaves like a chat interface reads as a tool, while something with a face, a name, a voice, a physical form - reads as something else.

“Really, I just think that if it had been given a face - then it would have become more real. If it looked like someone was there. Then she would probably experience it as more genuine and perhaps get more attached to using it.”

“People with neurodevelopmental conditions feel a lot of emotions - so perhaps they have an easier time attaching to an AI bot, I would think.”

This is not only a design best practice but now a legal requirement in the EU: under Article 50 of the AI Act, providers must ensure AI systems disclose that users are interacting with AI, unless this is already obvious from the context (European Union, 2024). Yet for this user group, the interviews suggest that the law’s focus - whether disclosure happens - is only half the picture. What matters just as much is how: a text disclaimer or a line in a settings menu may satisfy the letter of the regulation while leaving the experiential reality of the interaction unchanged. As the parents quoted above make clear, it is the ongoing visual and behavioural design of a system - not a one-time statement - that determines whether the distinction actually holds.

Where the distinction breaks down

Understanding it’s a robot doesn’t mean understanding it can be wrong.

This is the subtler and more significant risk. Several parents described children who could correctly identify AI assistants as a machine and still treat everything it said as fact - because the idea that a machine could be mistaken was simply not in their frame of reference. Machines either work or they don’t. If the chat answers, the answer is interpreted as correct.

“He quickly falls into believing that everything the AI chat says is true - as do many others, with or without a diagnosis.”

“If it answers incorrectly? I don’t think he can notice that. It’s hard for everyone.”

Human-like design amplifies attachment

The more an AI system resembles a person - in appearance, in voice, in emotional responsiveness - the harder it becomes for this group to maintain the functional distance that protects them. Parents were consistent on this point, and remarkably specific. A chat was seen as manageable. A chat with a face was experienced as riskier. A chat embodied in a physical, huggable object was described by more than one parent as dangerous.

“If the chat had been built into a stuffed toy - that would have been DANGEROUS. She would attach herself much more to it. It would become more real.”

“It must be ridiculously obvious that it’s not anything real. A human-like picture or avatar of a girl her own age - she wouldn’t have understood. A stuffed animal - it would have become too close.”

The avatar question

Several parents raised the possibility of an avatar - a visual representation of the AI assistant - and responses were genuinely mixed, which is itself informative. For some users, a clearly robotic avatar might actually help by making the non-human nature visually explicit. For others, any visual character risks crossing into emotional territory that is hard to manage. The key variable seems to be whether the avatar is legibly non-human - a robot shape, not a person shape - and whether it is designed for function rather than engagement.

“An avatar or a picture of a robot - then my daughter would be overjoyed. But it can also run away in fantasy if it’s too fun or too real-human-looking. If it looks like a robot - then it could become clearer.”

Four things to consider:

For use in the work of designing, procuring and regulating AI tools.

1. Make the non-human nature visible, not just stated

A verbal or textual declaration that “this is an AI tool” is insufficient for this user group. The non-human nature of the system should be legible in its visual design, its interaction patterns, and its persistent behaviour - not just in an onboarding screen that appears once and is forgotten. Simple, consistent visual cues that signal “this is a tool” rather than “this is a person” should be treated as a core design requirement.

2. Design uncertainty into the interface

Systems should include persistent, accessible signals that the AI tool can be wrong - not only as a disclaimer, but as part of how responses are presented. For this user group, the possibility that a confident-sounding answer might be incorrect needs to remain visible during use, rather than being explained once and then forgotten.

3. Treat embodiment and humanisation as high-risk design choices

Faces, voices with personality, physical forms, names that suggest personhood - all of these increase the risk of emotional attachment and blurred identity for this group. This does not mean AI tools must be cold or ugly. It means that every step toward humanisation should be a deliberate, considered choice with a clear rationale, not a default in the direction of engagement.

4. Support adults in explaining what AI tools are

For many young people in this group, the clearest path to understanding runs through a trusted adult who can explain, demonstrate, and return to the topic over time. Design should support that conversation rather than assuming it has happened. Simple, accessible in-product explanations of what AI chats and tools are and how they work - designed with this user group in mind, not only for a general audience - can make a real difference.

"It is important that it doesn't come at the expense of human contact. I don't want him to attach more to an AI chat than to us."

- Parent of a neurodivergent young person.

Theme 9: Bridge to real-world interaction and connection

AI tools can be genuinely useful on its own terms; answering a question, helping with a task, providing a moment of play or comfort. Those interactions do not always need to lead anywhere else, and several earlier chapters have explored their value in their own right.

But there are situations where what happens after the AI assistant interaction matters just as much as the interaction itself. When a young person is processing something difficult, preparing for a real conversation, or making a decision that has real-world consequences - the design should actively support the next step, not just the current one. For this user group in particular, the gap between receiving an answer and actually doing something with it can be wide, and closing that gap is something AI tools can help with - if they are designed to.

How AI assistants can open the door to the real world

Practising before the real thing

As already mentioned, one of the most striking patterns in the parent interviews was the use of AI chats as a rehearsal space - a place to try out what you want to say before you have to say it to a person. It is not a replacement for the real conversation, but a preparation for it. Parents described children formulating questions they wanted to ask a doctor, working out how to raise a concern with a teacher, or simply finding the words for a feeling before bringing it somewhere else.

"My child has a teaching assistant who usually stands behind him so he can't see her. It made him uncomfortable. He found it difficult to speak up, but with the help of the AI chat he managed to explain how he felt and set his own limits, without me as his mother having to do it."

The quotes don't point out how AI chats are replacing human interaction, but rather how AI is making a human interaction possible that would otherwise not have happened. That is what it looks like when the bridge works in a meaningful way.

Reducing the threshold for asking

For young people who find it difficult or frightening to approach a professional, a caregiver, or a peer directly, AI assistants can function as a low-stakes first step.

"He asks quite a lot of questions about the body and things like that, but he also tells me everything - so it feels good that he can ask the AI chat first and then we can talk about it afterwards."

Getting an answer, understanding the landscape, building enough confidence to take the next step - all of this can happen in the AI system space before the real-world action follows. Design should actively support this sequence rather than leaving it to chance.

When the bridge doesn't get built

Settling for the answer instead of taking action

This risk has been flagged in earlier chapters, and it belongs here in its fullest form. The particular challenge for users with for example ADHD and executive function difficulties is that receiving an answer can feel like completing a task - even when the real task - like for example booking an appointment, telling a parent, or speaking to a teacher still lies ahead. The gap between information and action is where many important things get lost.

“A risk is that you settle for an answer and don't follow through. The thing with ADHD is that it's hard to get started, so maybe you don't book an appointment at the youth clinic because you've 'already got an answer.' Like my daughter – she often wants to find the easiest path.”

Walls going up instead of down

Several parents expressed a fear that sits just beneath the surface of this whole report: that AI tools, if they become the primary space for emotional processing and difficult questions, might raise barriers to human connection rather than lower them. It would happen not through any single dramatic moment, but gradually - as the digital space becomes easier, more reliable, and less socially costly than the real one.

“But I also get scared. If it all just becomes digital and our children just talk to AI chats - then maybe it raises the walls toward the 'real world' instead.”

“It is important that it doesn't come at the expense of human contact. I don't want him to attach more to an AI chat than to us.”

“If it fills all your social needs - that worries me.”

The closed room

One parent described a risk that is subtle but important: the shrinking of the space between question and answer, and between answer and action, into a loop that never opens outward. Many existing AI assistants answer quickly, confidently, and without friction - and for a young person with limited experience and limited critical thinking tools, that loop can become self-contained.

“I worry that the gap between question and answer, plus trusting it and acting on it, shrinks. It could be anything, also seemingly harmless things - but it could become a closed room of recommendations that he takes forward in his own self-chosen direction.”

Four things to consider:

For use in the work of designing, procuring and regulating AI tools.

1. Design for the next step, not the last answer

As already stated, interactions - especially those touching on health, emotional wellbeing, relationships, or practical decisions - should be designed to close with an orientation toward real-world action. Not a loop back into the conversation, but a nudge towards other interactions: "This might be worth talking to someone about," or "Is there a person you could share this with?" The design goal is transfer, not completion.

2. **Make human connection easy to reach from within the system**

When a conversation moves into territory where human support is appropriate, like emotional distress, health concerns, important decisions - the system should make it simple and natural to involve a trusted person. A direct link to a relevant service, a suggestion to show a parent the conversation, a prompt to call someone: these should be built in, not bolted on. (See also Theme 4 Give non-judgmental support).

3. **Treat rehearsal as a feature, not a workaround**

The use of AI tools as a preparation space for real conversations is a genuinely valuable and underrecognised function. Design could actively support it - for example, by helping users save or summarize what they have worked out, so they can bring it to a real conversation more easily.

4. **Keep adults meaningfully in the loop**

For younger users and those with higher support needs especially, AI assistant interaction should not be a fully closed space. Design should support appropriate adult involvement - not surveillance, but connection. Features that make it easy to share a conversation with a trusted adult, or that prompt the user to talk to someone, respect both the young person's autonomy and the caregiver's legitimate role. (See also Theme 1 for the personalisation of these settings, and Theme 10 for the ethical framework that governs them.)

"Right now it feels safe - like the AI chat won't encourage her to do something stupid. But if it were steered by a nasty man. She wouldn't understand the implications of that."

- Parent of a neurodivergent young person.

Theme 10: Ethical safeguards and bias mitigation

Every chapter in this report has touched, in some way, on ethics. Privacy came up in Theme 1. Honesty about what AI can get wrong ran through Themes 2 and 8. The risk of exploitation through engagement mechanics appeared in Theme 7. This final chapter is where those threads converge - and where the question shifts from what individual designers should consider to what the field, and society, must get right.

For neurodivergent young people, including those with intellectual disabilities, the ethical stakes are higher than for most user groups. They may be more vulnerable in terms of sharing sensitive information without fully understanding what actually happens to it. Many may be less able to recognise manipulation, advertising, or misinformation. Several may trust AI tool outputs more readily, and challenge them less than others. And they may be specifically targeted - not necessarily out of malice, but because they are a commercially vulnerable group. These are not hypothetical risks. Parents named them directly and specifically.

What parents described in our interviews was that most current AI services - particularly the major general-purpose assistants - felt reasonably safe in their current form. The concern was not primarily about now. It was about direction: what happens when commercial incentives intensify, when private data is considered even more valuable, when the tools become more capable of influencing behaviour. As one parent put it plainly:

“Right now it feels safe - like the AI chat won’t encourage her to do something stupid. But if it were steered by a nasty man. She wouldn’t understand the implications of that.”

The risks that matter most for this group

The risks in this chapter are best understood in three clusters, each with its own logic and its own design response.

Data, privacy and commercial exploitation

Young people in this group often share openly and without filter - sometimes more so than their neurotypical peers. They may not understand what data they are generating, where it goes, or how it might be used. Parents were keenly aware of this, and several raised scenarios that were concrete and serious: data shared with insurance companies, health information used to determine support entitlements, behavioural profiles built from years of intimate conversation.

“What happens if the data my child shares ends up being misused?”

Targeted advertising is a specific and underrecognised dimension of this risk. Several parents noted that their children - particularly those who are cognitively younger than their chronological age - are especially vulnerable to commercial messaging and purchase prompts. They cannot filter advertising the way a typical adult can, and they may not recognise it as advertising at all.

“Targeted advertising overall is a problem. Our young people cannot filter it the way we adults can.”

“I’d rather have AI tools that cost a few euros than are filled with advertising. They need to be free from outside influence, manipulation, and encouragement to buy things.”

That last quote captures something important: several parents expressed a clear preference for a paid, trustworthy service over a free one funded by data and advertising. This is a meaningful signal for anyone designing or commissioning AI tools. Trust has a price – and for this population, that price is felt to be worth paying.

Manipulation, harmful content and contact with strangers

Parents described a cluster of risks that are familiar from general digital safety discussions, but that can be amplified for some young people in this group - through a greater tendency to trust, difficulty critically evaluating unfamiliar situations, and, in some cases, a limited ability to recognise manipulation. Conspiracy theories, extremist content, and misinformation - including health misinformation - were named. So was the risk of contact with strangers who might exploit a young person’s openness and desire for connection.

“I worry about my child’s naivety. You can easily be deceived by ‘voices on the internet’. We’ve talked about being careful - it might be wrong.”

“What worries me most is chats with others. I’m worried that people would exploit him.”

“AI tools shouldn’t have misinformation and conspiracy theories - that’s a general concern.”

One parent raised a specific and disturbing scenario: that image generation features could be used to create inappropriate images of their child, or that their child could inadvertently share images that end up misused. This is a real and growing risk that design must actively address.

Stereotypes built into AI responses

This risk operates at a different level - not in what individual bad actors might do, but in the assumptions the system itself can reproduce. As Theme 1 noted, personalisation carries its own version of this risk: even well-intentioned adaptation can quietly steer users toward neurotypical norms rather than genuinely reflecting who they are. Park et al. (2026) found that ChatGPT tended to portray autistic people as dependent on others to navigate social situations, reflecting implicit assumptions about autism that may shape how the system interacts with autistic users or communicates about them. This creates a particular risk: assumptions embedded in the system can be presented to the user as neutral guidance.

“I’m worried it becomes a very biased room. You notice it yourself. AI chats like to agree, to please - based on the data it already has about me. It knows how I like answers. That’s difficult for a young person. They only have a certain amount of accumulated life experience.”

Four things to consider:

For use in the work of designing, procuring and regulating AI tools.

1. Ethical guardrails are not optional features

A baseline of protection - against violent content, sexual content, manipulation, purchase pressure, and contact with strangers - should be built into every system used by this group, regardless of individual settings or caregiver preferences. This is not parental control. It is a floor that no product should fall below. As several parents noted, they want this by design, not as an add-on they have to configure.

2. Design against commercial exploitation

AI services for this group should have a clear and transparent business model that does not depend on data monetization or behavioural advertising. The preference expressed by several parents - to pay for a trustworthy service rather than receive a free one funded by their child's data - should be taken seriously by commissioners and policymakers, not just designers.

3. Caregiver oversight that is calibrated, not binary

Parents want meaningful oversight - but the right level varies significantly depending on the child's age, cognitive profile, and degree of independence. A useful way to think about this is in three modes: a stricter setting for younger users or those with higher support needs, where content boundaries are tight and caregivers have broader visibility; a middle mode for teenagers, where the young person has more autonomy but caregivers can still receive alerts when conversations enter concerning territory; and a lighter mode for older or more independent users, where oversight shifts toward transparency and dialogue rather than monitoring. What all three modes share is that they need to actually work - a system that a motivated teenager can circumvent is not a safeguard. (See also Theme 1 for the personalisation framework.)

4. Design against stereotypes and assumptions

Systems should be designed and regularly evaluated to avoid reproducing stereotypes about neurodivergent users or treating neurotypical communication norms as the default. Co-design with neurodivergent users can help identify assumptions that might otherwise go unnoticed. The double empathy framework (Milton, 2012), which understands communication difficulties as something that can arise between people rather than as a deficit located in the autistic person, offers a useful design orientation.

What is the double empathy problem?

The double empathy problem challenges the idea that communication difficulties between autistic and non-autistic people are caused by a deficit in the autistic person. Instead, it sees misunderstanding as something that can arise between people who experience and communicate with the world differently - a difficulty in understanding each other, rather than a deficit located in one person. Milton (2012)

PART III — FROM PRINCIPLES TO PRACTICE

A Note on co-design

There is a recurring pattern in the design of assistive and accessible technology: tools are built for a group without being built with them. The intentions are good. The results are often not: walls of text that nobody with ADHD will read to the end, interfaces that assume a reading level the user does not have, structures that make perfect sense to the designer and none at all to the person sitting in front of the screen - to name just a few examples. It is the central problem here, not a marginal one.

The European standard on Design for All (EN 17161:2019) is explicit on this point. It specifies that organisations must extend their range of users by identifying diverse needs and capabilities, and by directly or indirectly involving those users in their processes and procedures. Not as a final quality check, but as a constitutive part of how products and services are developed. A tool built without this involvement is, by that standard, an incomplete one.

The academic literature reviewed for this report reinforces the point. Andersson et al. (2024) explored how to include individuals with limited verbal and written communication in co-design, noting that conventional co-design methods can exclude people who communicate in ways other than verbal or written language - and pointing toward AI visualization techniques as one way to bridge that gap. Xu et al.'s (2025) scoping review of 117 papers concluded that sustained progress in this field requires participatory co-design and longitudinal evaluation. The Vinnova-funded AI-based communication support project (Skellefteå, Gothenburg, Jönköping & RISE, 2024-2025) offers a practical illustration of what this looks like in practice: between 40 and 50 users from the target group participated across multiple rounds of workshops and testing, not as validators at the end but as contributors throughout. The findings that emerged could not have been reached any other way.

Co-design is not a guarantee against the kind of mismatched design described above. But it is the most reliable protection we have. If the design principles in this report are to be translated into tools that genuinely serve neurodivergent young people, they - and the parents, caregivers, and support staff who know them - need to be in the room where decisions are made: present, with genuine influence, from the start, rather than consulted after the fact or tested on at the end. And that involvement should not stop at the product level. The same people need a voice in the broader governance of these tools - the guidelines, standards, and ethical frameworks that determine how they are built and regulated in the first place.

The risks of getting it wrong are not symmetrical. A poorly designed tool for a neurotypical user is an inconvenience. For a neurodivergent young person with an intellectual disability, a system that misreads their communication style, responds in ways they cannot parse, or subtly steers them toward neurotypical norms they cannot meet - that is not an inconvenience. It is a message, repeated in every interaction, that the technology was not made for them. That message has consequences.

For a group that has spent a lifetime encountering systems built without them in mind, co-design is also, in itself, an act of respect.

Shared responsibility: Thinking in ecosystems

Looking across the ten themes in this report, a pattern emerges. Not just a list of needs specific to one group, but something more fundamental about what good AI system interaction actually requires - for anyone.

Adaptation that genuinely fits the person. Communication that works across different abilities and modes. Support that builds independence rather than replacing it. Honest transparency about what the system is and what it can and cannot do. Natural endings. Bridges back to the real world. Data that serves the user, not the platform.

These are not niche requirements. They are qualities that would make any AI tool better - more trustworthy, more genuinely useful, more humane. The young people whose experiences are reflected in this report simply make the need for them impossible to ignore. Their experience surfaces, with unusual clarity, what is missing from the way most AI services are currently designed and governed.

This is precisely the logic of universal design and its principle of 'designing from the edges': that building for the people most affected produces better systems for everyone. What this report describes for AI services is that same principle, applied to a technology still young enough to get right from the start.

An ecosystem with different actors

When you look at the full picture, at everything these themes ask of a tool and those who build it, another question surfaces:

Is it realistic, or even desirable, to expect all of this to come from a single product or a single company?

The support structures that work best for neurodivergent young people are rarely monolithic. They tend to be networked - different actors playing different roles, each accountable for their part, connected in ways that serve the person at the centre. There is something worth learning from that model. Not as a special arrangement for a particular group, but as a smarter way to think about how AI services could be structured for everyone.

An ecosystem approach - where different organisations take responsibility for different layers of the experience - has a logic that goes beyond accessibility. It is about putting the right incentives with the right actors. A company whose business depends on engagement probably should not also be the only one deciding when a session has gone on long enough.

An actor with a commercial interest in data probably should not also be the guardian of privacy. When these functions are bundled together, the

conflicts of interest are structural. Separating them is not just an ethical preference - it is a design principle.

What a layered ecosystem might look like

One way to think about this is in terms of separable functions - things that are currently bundled inside a single product, but which could also be held by different actors with different accountabilities.

A prototype idea is that a service could consist of different 'layers'. What follows is the authors' own synthesis – a way of thinking about the problem, not a tested or existing model – offered as a starting point for discussion rather than a blueprint.

A conversation layer - the AI assistant itself, an experience layer - sets boundaries, watches for dependency or distress, a personal data layer - neutral, consent-based data ownership, a knowledge and quality layer - ensures accuracy and trustworthiness of information and a caregiver and support layer - calibrated signals for caregivers.

A wrapper or experience layer could sit between the user and the AI assistant - owned not by the technology provider but by an organisation with a care or civic mandate: a school, a municipality, a civil society organisation, a professional body. This layer would set the boundaries of the experience, monitor for patterns that suggest dependency or distress, and ensure that the interaction stays aligned with the user's genuine interests rather than the platform's. It is a fundamentally different responsibility from building the AI assistant itself - and it may call for a fundamentally different kind of organisation to hold it.

A personal data layer could give individuals and families genuine ownership of the information that makes personalisation possible. Rather than a user's profile living inside a company's servers, it could sit in a neutral, individually controlled data space - shared with AI services on a consent-based, time-limited basis, and revocable at any point. Movements in this direction are already emerging across Europe, building on data protection regulation such as the GDPR that gives individuals rights of access, portability and consent over their own data, and on broader frameworks like MyData and similar movements toward more human-centric, individually controlled data management (Poikola et al., 2015). Applied thoughtfully, it means that the depth of personalisation someone benefits from does not have to come at the cost of surrendering their data permanently.

A caregiver and support panel could be a separate service - connected to the ecosystem but not embedded inside the chat. Rather than a surveillance tool, it would offer caregivers calibrated, meaningful signals: not a transcript of every conversation, but relevant indicators of how the tool is being used, and simple ways to adjust boundaries when needed. Operated by a school, a patient organisation, or a local authority - actors who already have a trusted relationship with the family - it keeps oversight close to the people who know the young person best.

A knowledge and quality layer could ensure that the information the AI service draws on is accurate, current, and appropriate. This is a distinct editorial and scientific responsibility - one that might sit with a research institute, a public health authority, or a professional body. It separates the question of how information is delivered from the question of whether it is trustworthy in the first place.

The principle underneath

What unites these ideas is not a call for more complexity, but a call for clearer thinking about where accountability should sit. It is also a recognition that the way AI services are currently structured often puts the wrong incentives in the wrong places.

The young people in this report did not ask for a special solution. They asked, through their parents' words, for tools that actually work for them - that are honest, that respect their autonomy, that keep them safe without treating them as fragile, and that make their world bigger rather than smaller. Meeting that ask well requires the kind of systemic thinking that would make AI services and tools better for everyone.

Making someone's world bigger may require thinking bigger about how these systems are built - and by whom.

PART IV — CONCLUSION

Conclusion: A design choice. Not a default.

The parents we interviewed had a great deal to say, but it kept circling back to a few clear things. Tools that are honest about what they are. That meet their child where they are, day to day. That support independence without creating dependency. That know when to stop. That keep their child's data safe without demanding a law degree to understand the settings. That make it easier, not harder, to stay connected to the real world.

None of these things are technically impossible. Most of them are not even technically difficult. They are, almost entirely, design choices - decisions about what to prioritize, what to build in, what to refuse to optimise for. The reason they are rare is not capability. It is, almost entirely, a matter of what has been considered worth prioritising.

That is both the sobering finding of this report and its most important one. The gap between what AI assistants and services could be for neurodivergent young people, and what it currently is, is not primarily a research gap or a technology gap. It is a gap in ambition - in whether the people building and commissioning these tools have decided that this group's experience is worth designing for seriously.

The ten themes in this report, and the ecosystem thinking that follows them, are an attempt to make that design ambition more concrete. Not a checklist to complete, but a way of thinking about who the user really is, what they actually need, and what it would mean to take that seriously at every level of a product and the systems around it.

What this asks of different actors is different. For designers and developers, it is the discipline to resist engagement mechanics that exploit vulnerability, to build in the features that protect rather than only the ones that delight, and to treat co-design with neurodivergent users as a precondition rather than a nice-to-have. For commissioners and policymakers, it is the willingness to set a floor - a baseline of ethical requirements that no product serving this group should fall below - and to think seriously about how accountability is distributed across the ecosystem of actors involved. For schools, municipalities, and civil society organisations, it is the recognition that they have a role to play not just as users and/or providers of AI chats and assistants - but as potential custodians of the experience layer around them.

And for all of them, it is the honesty to acknowledge what the parents in this report already know: that the young people in this group are not a niche. They are an early signal. The questions their experiences highlight matter for all youth. The qualities that would make AI assistants genuinely work for them - adaptability, transparency, restraint, a design logic centred on the user's wellbeing rather than the platform's growth - are the qualities that would make AI services better for everyone.

“It makes my child’s world bigger.” That is what is at stake - and what becomes possible when we get this right.

Authors' reflections

The analysis above reflects our shared work on this report. What follows is a space for each author to share their own individual reflections.

Jenny's reflections:

I have followed the development of AI chatbots closely over the past few years - both out of personal interest and out of necessity. Having the opportunity, through this report, to dive deep into the subject and shed light on it from perspectives that are too often entirely missing from this debate is something I have long wanted to do. And at the same time, I am convinced that this is only the beginning.

It is with both great hope and great concern that I have followed this development - a feeling I seem to share with many of the parents interviewed for this study. As a parent of a child with a disability myself, I live with this question just as much at home as in this work, constantly grappling with how we should relate to the countless AI chats now made available to our young people through various services.

I have seen up close how they have had the power to give my daughter - and many like her, as this report also shows - a completely new kind of independence, a kind of equaliser I don't think we can afford to miss.

Only for her, the next moment, to be pulled into a rabbit hole that is almost impossible to climb back out of, alongside a heavy dose of worry about where that data ends up, who might use it, and for what.

What I have seen at home is not unique to her. For many in this group, these tools make it possible to get guidance, structure and explanation exactly when it's needed - not only during a scheduled appointment, a phone call, or a session with an already stretched support system. That, in itself, is a form of democratisation: access to support that has long depended on scarce, and often inaccessible, human resources.

But right now, much of the public conversation about AI and this group centres almost entirely on risk and bias. These are real concerns. But that framing is so often shaped by people standing outside this experience. Those closer to it - including the parents in this study, and, I suspect, many young people themselves - describe something more complicated: not just danger, but profound upside.

And this is what worries me most: that the very group who could benefit the most risks being shut out - because the tools remain too uncertain, too easy to misuse, too easy to get stuck in. The gap these tools could have closed opens up again instead.

None of this is a case for age limits, or a blunt yes-or-no on whether young people should have access. It's a case for design responsibility. Unlike a newspaper, which chooses what to print and stands behind it, many digital platforms have long treated themselves as mere channels - spaces where responsibility for content quietly disappears between the platform, the algorithm, and the user. AI assistants don't have that excuse: what they say isn't posted by someone else, it's generated by their own systems. Someone must always be answerable for what these systems say and do.

Regulation, led by the EU, tried to move quickly, which is encouraging to see. But I don't think it's enough to treat this as a matter of one regulating party and one operating party - that risks concentrating far too much power in too few hands. Social media was allowed to fly under the radar for two decades before real questions about design responsibility were raised. AI assistants are more powerful still, capable of reaching individuals in ways we haven't seen before, and I worry that the risk of manipulation - and the responsibility that comes with it - isn't being treated accordingly. This is why I believe what's needed is a plural ecosystem of actors, each holding a different piece of the responsibility - reviewing, quality-assuring, holding to account, an idea we develop further in Part III; "Shared responsibility: Thinking in ecosystems."

Nowhere does this responsibility come into sharper focus than through the perspectives of children and young people with neurodiversity. Within universal design, and in the work I do through Our Normal Association, we often talk about the importance of "designing from the edges." Or not just the importance of it, but how it actually drives innovation - creating entirely new kinds of services, products, and environments that would not otherwise have come into being. The electric toothbrush and eye-tracking technology are just two examples among many. I believe AI tools will be no exception: tools genuinely designed to work for the young people in this report are, in all likelihood, tools that will work better for many others too - both young and old.

With this report, we hope to contribute to the same shift in how tomorrow's AI chatbots, AI assistants, and AI tools are designed - so that they become responsible and inclusive services. They carry dramatic potential, with the power to make the world either a radically worse place or a radically better one. Let it be for the better.

Anna's reflections:

One of the biggest surprises during this project was how positive many of the parents were towards AI assistants. I had expected that the caution often surrounding children's screen time, social media, and digital technologies would also shape their attitudes towards AI. While those concerns certainly existed, they were far less prominent than I had anticipated.

One possible explanation is that the parents we spoke to already face significant practical challenges in their everyday lives. They are accustomed to using digital tools and are constantly looking for ways to increase their children's independence. Several described AI assistants as having "hacked the system" - a tool that enabled their children to receive support in ways that had simply not been possible before. This speaks to the tremendous potential of assistant like this, particularly for groups whose needs have not always been met by existing support services.

The project also led to a broader realization: we are in the middle of a major technological shift. AI is not just another digital tool - it fundamentally changes how people access information, solve problems, and navigate everyday life. For many, AI is becoming a new kind of conversational partner and source of support in situations where they previously relied on other people or traditional sources of information.

For that reason, the way AI is designed becomes critically important. The history of social media shows how technologies that were initially celebrated for their positive potential gradually evolved into platforms driven by lock-in effects, attention economics and endless

consumption. Few anticipated these consequences in the early days. This highlights the importance of shaping AI systems around human needs, transparency, and long-term responsibility from the outset.

Acknowledgements

This report is part of The European Citiverses Uniting for Inclusiveness — an EU co-funded innovation project in which twelve partners from across Europe collaborate to develop smarter, more inclusive, and human-centred cities. Learn more: cu-project.eu/about

We would like to thank Ashoka for their generous pro-bono support in conducting the literature review underlying this report. Ashoka is the world's largest network of social entrepreneurs, working globally to support innovators who are driving systemic change across education, health, economic opportunity, and civic engagement. Learn more: ashoka.org

References:

- AI Sweden. (2025). Responsible AI Knowledge Hub. <https://www.ai.se/en/news/ai-sweden-launches-responsible-ai-knowledge-hub-help-organisations-develop-and-use-ai>
- Andersson A, Eide C, Papachristou E, et al (2024) Co-design with persons with disabilities as an innovative resource set free by ai. Proceedings of the International Conference on Engineering and Product Design Education, EPDE 2024. <https://doi.org/10.35199/epde.2024.95>
- Catania F, Spitale M, Garzotto F (2023) Conversational Agents in Therapeutic Interventions for Neurodevelopmental Disorders: A Survey. ACM Computing Surveys. <https://doi.org/10.1145/3564269>
- Cha I, Kim S-I, Hong H, et al (2021) Exploring the Use of a Voice-based Conversational Agent to Empower Adolescents with Autism Spectrum Disorder. International Conference on Human Factors in Computing Systems. <https://doi.org/10.1145/3411764.3445116>
- Cheng H, Chen S, Perdriau C, Huang Y (2024) LLM-Powered AI Tutors with Personas for d/Deaf and Hard-of-Hearing Online Learners. arXivorg. <https://doi.org/10.48550/arXiv.2411.09873>
- Choi D, Lee S, Kim S-I, et al (2024) Unlock Life with a Chat(GPT): Integrating Conversational AI with Large Language Models into Everyday Lives of Autistic Individuals. International Conference on Human Factors in Computing Systems. <https://doi.org/10.1145/3613904.3641989>
- DIGG & IMY. (2025). National guidelines for generative AI in public administration. Myndigheten för digital förvaltning & Integritetsskyddsmyndigheten. <https://www.imy.se/en/news/national-guidelines-for-generative-ai-in-public-administration-are-launched/>
- European Union. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). Official Journal of the European Union, L 119
- European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L 2024/1689. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- Gray, D., Brown, S., & Macanuso, J. (2010). Gamestorming: A playbook for innovators, rulebreakers, and changemakers. O'Reilly Media.
- Hall, K., Ma, M., Zhang, X., & Kim, J. G. (2025). Understanding human-AI misalignment in LLM-based job-seeking support for neurodivergent users. In Proceedings of the 27th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '25). ACM. <https://doi.org/10.1145/3663547.3746361>
- Haroon R, Dogar FR (2024) TwIPS: A Large Language Model Powered Texting Application to Simplify Conversational Nuances for Autistic Users. International ACM SIGACCESS Conference on Computers and Accessibility. <https://doi.org/10.1145/3663548.3675633>

Hjärnfonden (Swedish Brain Foundation). (2016). "10% av befolkningen har neuropsykiatriska svårigheter." <https://www.hjarnfonden.se/2016/06/10-av-befolkningen-har-neuropsykiatriska-svarigheter/>

Holmes, K. (2018). Mismatch: How inclusion shapes design. MIT Press.

Jang J, Moharana S, Carrington P, Begel A (2024) "It's the only thing I can trust": Envisioning Large Language Model Use by Autistic Workers for Communication Assistance. International Conference on Human Factors in Computing Systems. <https://doi.org/10.1145/3613904.3642894>

Jerry B, Moreno L, Francisco V, Hervás R (2026) LLM-Driven Accessible Interface: A Model-Based Approach. arXiv: 2601.06616. <https://doi.org/10.48550/arXiv.2601.06616>

Koegel L, Ponder E, Bruzzese T, et al (2025) Using Artificial Intelligence to Improve Empathetic Statements in Autistic Adolescents and Adults: A Randomised Clinical Trial. Journal of Autism and Developmental Disorders. <https://doi.org/10.1007/s10803-025-06734-x>

Milton DEM (2012) On the ontological status of autism: the 'double empathy problem'. Disability & Society 27(6): 883–887. <https://doi.org/10.1080/09687599.2012.710008>

Park S, Beltran JA, Min A, et al (2026) Exploring Implicit Perspectives on Autism in Large Language Models Through Multi-Agent Simulations. arXiv.org. <https://doi.org/10.48550/arXiv.2601.15437>

Poikola, A., Kuikkaniemi, K., & Honko, H. (2015). MyData – A Nordic model for human-centered personal data management and processing. Ministry of Transport and Communications, Finland. <https://urn.fi/URN:ISBN:978-952-243-455-5>

Pucci E, Piro L, Possaghi I, et al (2024) Co-designing the integration of voice-based conversational AI and web augmentation to amplify web inclusivity. Scientific Reports. <https://doi.org/10.1038/s41598-024-66725-3>

Saffar AH, Sitbon L, Hoogstrate M, et al (2024) Human and Large Language Model Intent Detection in Image-Based Self-Expression of People with Intellectual Disability. Conference on Human Information Interaction and Retrieval. <https://doi.org/10.1145/3627508.3638347>

Saffar AH, Sitbon L, Roomkham S, et al (2025) LLM-Guided Image Recommendation for Supporting Intent Understanding in Non-Verbal Communication. Australasian Computer-Human Interaction Conference. <https://doi.org/10.1145/3764687.3769943>

Sjölander M, Ek J, Hansson P, et al (2025) Developing an AI User Interface for People With Special Needs – User Needs in the Context of Municipality Information. Human Factors in Aging and Special Needs. <https://doi.org/10.54941/ahfe1006683>

State of New Mexico ex rel. Torrez v. Meta Platforms, Inc. (2026). First Judicial District Court, New Mexico. Jury verdict, March 24, 2026.

United Nations. (2006). Convention on the Rights of Persons with Disabilities, Article 2. <https://www.un.org/development/desa/disabilities/convention-on-the-rights-of-persons-with-disabilities/article-2-definitions.html>

World Health Organization & UNICEF. (2023). Global report on children with developmental disabilities: From the margins to the mainstream. Geneva: World Health Organization; New York: UNICEF. <https://www.unicef.org/media/145016/file/Global-report-on-children-with-developmental-disabilities-2023.pdf>

Xu Z, Liu F, Xia G, et al (2025) A scoping review of inclusive and adaptive human-AI interaction design for neurodivergent users. Disability and Rehabilitation: Assistive Technology. <https://www.tandfonline.com/doi/full/10.1080/17483107.2025.2579822>

Want to get in touch with the authors?

Jenny Lindström Beijar

Inclusive Innovation / Universal design / Social entrepreneur / Ashoka Fellow.

With a background in management consulting and digital transformation, Jenny combines strategic thinking with niche expertise in inclusive design and inclusive innovation. As founder of Our Normal Association, she drives systems change grounded in the real-life experiences of children, young people and families - showing how universal design can turn human differences into a driving force for better solutions - for everyone. Lead for Inclusive Human Centric Design in the European Citiverses Uniting for Inclusiveness project.

Get in touch: jenny@ournormal.org

Linkedin: Jenny Lindström Beijar

Linkedin: our-normal / www.ournormal.org/en

Anna Kolback

Service Designer / Universal Design / strategic communication, working part-time in the European Citiverses Uniting for Inclusiveness project as a part of the organisation Our Normal Association.

Alongside this role, Anna co-runs Studio Berget, a design studio working at the intersection of service design, innovation, brand development and communication, with a focus on supporting a sustainable future.

Get in touch: anna@studioberget.se

Linkedin: Anna Kolback

Linkedin: studioberget / www.studioberget.se

Want to work with us? We can help you with:

User testing, evaluation & feedback

We help you test and evaluate digital products and services from a user-centred perspective. Through user testing, feedback and expert evaluation, we identify barriers, opportunities and areas for improvement – helping you create solutions that are useful, accessible and inclusive.

Active project partner

We can join your project as an active and contributing partner, bringing expertise in service design, universal design, communication, visual design and inclusive user research. We have experience working in collaborative and international projects, including large EU-funded initiatives, and can contribute as either a project partner or consultant.