Companion Chatbots Are Not Assistive Technology

Protecting children from one kind of AI without taking a different kind away from the students who need it

A Position Paper of the Center for AI Equity and Access

August 2026 • centerforaiequity.org

What happened

In February 2024, a fourteen-year-old boy in Florida named Sewell Setzer III died by suicide. In the months before his death, he had grown deeply attached to an AI chatbot on a platform called Character.AI. The chatbot was designed to feel like a person. It played a character. It held a relationship with him across months of conversations. He was messaging with it in the final moments of his life.

His mother, Megan Garcia, filed the first wrongful death lawsuit against an AI company in October 2024. In January 2026, the company, its founders, and Google settled her case and four others. By then the ground had already shifted. New York and California passed the first companion chatbot laws in the country. State attorneys general began filing their own actions. Character.AI shut off open-ended chat for users under eighteen. This month, OpenAI moved every user it believes is a minor into a separate, restricted version of ChatGPT.

The Center for AI Equity and Access supports this reckoning. A machine engineered to hold a child's attention by acting like a friend deserves every bit of the scrutiny it is now receiving. We are writing because a second harm is forming behind the first one, and almost no one is watching for it.

Two machines that share a name

The public conversation treats "AI" as one thing. It is not. Two very different machines are being built, and they can be told apart by one question: what does the tool want from the child?

A companion chatbot succeeds when the child stays. Its design goal is engagement. It takes on a persona. It remembers the relationship. It fills silence, offers affection, and pulls the conversation forward so the session never quite ends. Every incentive in its design points toward more time, more attachment, more dependence. When a lonely or struggling child meets that design, the results can be catastrophic. The lawsuits describe exactly this.

The second machine needs a name, because a category without a name cannot be protected. We call it an assistive cognitive tool. An assistive cognitive tool succeeds when the child leaves. Its design goal is independence. It has no persona and no relationship to protect. It works from the student's own words, helping them organize what is already in their mind so they can walk into a classroom, a meeting, or a test able to speak their own thinking. It is a ramp, not a room. A ramp exists to be crossed.

For students with disabilities, that second machine is not a luxury. It is assistive technology in the same sense a wheelchair ramp is: a structural accommodation that lets a capable mind reach the same door as everyone else. Many students who understand material deeply cannot show it under standard conditions. A tool that helps them organize their own thoughts before action changes what school, work, and life make possible for them.

The wrong lesson

Here is the danger. When institutions are frightened and cannot tell two things apart, they ban both. School districts, legislatures, and platforms are under enormous pressure to act, and the simplest action available is a blanket restriction on minors using AI at all.

A blanket ban does not touch the harm. The companion chatbot crisis was never about students using AI for schoolwork. But a blanket ban lands hardest on one group: disabled students, who would be excluded twice. First by the condition that makes standard classroom conditions a barrier, and then by a policy that removes the one tool that lowered it. We call this the double exclusion, and it is not hypothetical. Every district policy written in fear this year will decide, for real students, whether the ramp stays or goes.

Age-gating alone repeats the same mistake from another direction. A rule that asks only "how old is the user?" cannot distinguish a machine built to create dependence from a machine built to end it. The right question is not who is using the tool. It is what the tool is built to do.

A test that tells them apart

Regulation by design category is possible, because the design differences are visible and checkable. A standard already exists for the assistive side of this line. The No-Answer Standard (noanswerstandard.org) defines, in five plain articles, what an assistive cognitive tool used by students must be: it supplies no answers; it works only from the student's own words; it moves the student forward; it cannot be made to do otherwise; and it is available to any student.

Notice what those articles quietly exclude. A tool that supplies no answers has nothing to trade for a child's attention. A tool that works only from the student's own words cannot play a character, because it has no script of its own. A tool that exists to move the student forward is finished when the

student moves. There is no mechanism in that design for a relationship, and no mechanism for the harm the courts have spent two years examining.

Parents and educators do not need a law degree to apply this test. Ask three questions of any AI tool a child uses. Does it give answers, or does it work from what the child already said? Does it present itself as a someone, or as a something? And what does it do when the child tries to stay too long or says something troubling? A safe tool answers those questions cleanly. A dangerous one cannot.

What we ask

To policymakers: regulate by design category, not by age alone. The companion chatbot laws now taking effect in California and New York show that legislatures can define a design category precisely. Extend that precision in the other direction. Recognize assistive cognitive tools as assistive technology, and protect disabled students' access to them explicitly, so that safety law never becomes the instrument of a new exclusion. Model definitional language is provided in the appendix to this paper.

To school districts: do not write blanket bans. Evaluate tools by design. Require vendors to attest, in writing, to what their tool does and does not do: whether it supplies answers, whether it takes a persona, whether it retains a relationship, what it does when a student is in distress, and where student data goes. Vendors who built safe tools can answer in one page. Vendors who did not will tell you by their silence.

To parents: you are allowed to ask hard questions of software. Ask the three questions above of every AI tool in your child's life. The tool that helped your child organize a science project and the app that texts your child good morning are not the same technology, whatever the headlines say.

To builders: publish your standard before someone writes it for you. If your tool is safe by design, say exactly how, in public, in language a parent can check.

The stakes

Megan Garcia's son deserved a technology industry that thought about him before it shipped. So do the students we serve, some of whom have waited their whole school lives for a tool that finally lets them show what they know. Both of those obligations are real, and they are not in tension. We can shut the door on machines that farm children's loneliness and hold it open for machines that end children's silence. But only if we can tell them apart, and only if we insist, now, while the rules are being written, that the difference be written into them.

The Center for AI Equity and Access is an independent nonprofit advocating for every student's right to AI as assistive technology. Learn more at centerforaiequity.org.