



Insights from CES: Could Kid-Friendly Physical AI Agents Become the “Training Wheels” for AI Literacy?

Lyonne Zhu | Digital Safety Tech Policy Fellow | January, 2026

I just came back from CES in Las Vegas, and one product trend stood out to me more than dancing robots, Agetech, and faster chips: physical AI agents for kids. These are small, often cute, voice-enabled devices designed for educational use. They promise to answer questions, guide learning, tell stories, and support curiosity in a form that feels more like a learning companion than a screen-based app.

At first glance, this category might look like "edtech with a new shell." But I think it signals something bigger. Physical AI agents may become an early foundation for AI digital education, not just by teaching kids facts, but by teaching kids how to interact with AI itself. Unlike the "infinite scroll" of a tablet that excites dopamine receptors, a physical agent is **finite**. **It occupies space, it has a body, and it can be put away.** From a developmental perspective, this "embodied" AI offers a tangible boundary that abstract software does not.

That matters because in one CES session, "Youth Mental Health: Helping a Generation Thrive in a Digital World," speakers pointed out a core concern that kept surfacing: kids are entering online spaces with powerful tools and weak preparation. The session noted that schools generally do not have structured digital education that prepares youth to navigate technology safely and critically. The panel also acknowledged a hard constraint: building a full, modern digital education system requires public funding and long-term capacity, and government budgets often cannot meet the scale or speed of the need. A result of such funding constraints is that there are opportunities for those with deep pockets—such as tech companies—to step in and invest.

I agree with that solution, but I think we also need to think about what “investment” should look like in practice. The specific answers include curriculum grants, digital literacy nonprofits, and online modules, and the latter two include strong pushes for the safe use of AI. But families are still left with the everyday reality: kids may already have access to voice assistants, recommendation systems, and now increasingly general-purpose AI chatbots. Many

children will see AI first as a product with “consumer service interaction,” quoting from one of the speakers at CES.

That is where physical AI agents could play a real role. If these devices are designed specifically for kids with intentionally-limited scope, educational by default, and built with strong boundaries, they can function like training wheels for AI interaction. They can help children build habits that will matter later, when they graduate into much more open-ended systems like ChatGPT or Gemini. In other words, they can help young users **move from childhood curiosity to teen autonomy with fewer painful bumps.**

Crucially, this physical form factor allows for "positive friction." Kids do not learn digital literacy through a one-time lecture; they learn through repeated experiences and feedback. A screen-based AI is often optimized to be frictionless and instantly gratifying. A physical agent, however, can be designed to pause. It can "simulate confusion" to force a child to rephrase a rude or unclear command, teaching them that an AI is a tool that requires clear, critical input. This teaches the early skills of "prompt engineering" and critical thinking in a social, low-stakes sandbox.

This approach also has a security logic behind it. One of the biggest risks of "cloud-based" AI education is data privacy. Physical agents, however, are increasingly capable of **"on-device" processing** (Edge AI¹), meaning a child's conversation stays in the room instead of being uploaded to a server. This offers a concrete policy guardrail: a device that "lives" in the home, processes data locally, and is easier for parents to supervise because it is not an open portal to the entire internet.

Just as importantly, a physical form may shape behavior differently than an always-open chat app. A device that "lives" in a shared space at home can be easier for families to supervise and discuss. A child can treat it like a learning tool, not a private world. That subtle difference could matter for trust, safety, and the social development of kids who are growing up with increasingly personalized technology.

Still, the benefits depend on design choices. If the product simply becomes another attention machine, or another pseudo-friend optimized for engagement, then it will reproduce the same problems we already see in youth tech. But if it is intentionally built to support learning, curiosity, and critical thinking—and if it is transparent about what it is and what it cannot do—then it could support a healthier transition into the AI era.

¹ [What Is Edge AI? | IBM](#) Edge artificial intelligence refers to the deployment of AI algorithms and AI models directly on local edge devices such as sensors or Internet of Things (IoT) devices, which enables real-time data processing and analysis without constant reliance on cloud infrastructure.

This is why the most important question is not "Should kids use AI?" The real question is: **What kind of AI experience do we want kids to practice on first?** If their first exposure is an unrestricted, emotionally sticky, always-available conversational system, we should not be surprised when dependency, confusion, or harmful outputs show up. But if their first exposure is a carefully bounded educational agent, we might build a generation that approaches AI with more skill, more skepticism, and more agency.

In the short term, **physical AI agents are not a replacement for public digital education. They are not a substitute for parental involvement, school-based guidance, or safety regulation.** But they can be a **practical bridge.** And if tech companies are serious about youth safety and future user trust, investing in kid-centered physical agents paired with real AI literacy partnerships could be one of the most realistic steps they can take right now.

The transition is already happening. Kids are already growing up in an AI environment. The question is whether we are going to leave them alone in it, or give them a safer way to learn the rules before the world gets bigger. This leaves us with the most difficult work ahead: **defining the "shape" of this safety.** As these physical agents move from prototype to product, we must confront additional uncomfortable design and policy questions that go beyond simple parental controls:

- Should a child's AI agent look like a person or a pet, or does that anthropomorphism invite dangerous emotional dependency? Should we instead mandate "tool-like" industrial design, such as devices that look like radios or cubes, to visually reinforce that this is a utility, not a friend?
- If a child treats a physical agent as a confidant, who owns that "emotional data"? Do we need a "Right to be Forgotten" for hardware, where a parent can physically wipe the device's memory to ensure data privacy?
- Should regulators require these devices to have forced downtime where the AI refuses to engage to teach children that digital availability is not infinite?

After CES 2026, I am excited about such possibilities of the hardware holding a better digital future, and I look forward to answering these questions together.