



FAIR LABS

FAIR ARTIFICIAL INTELLIGENCE RESEARCH LABS

RESEARCH BRIEF ♦ CHILD SAFETY ♦ 2026

AI Companions and the Developing Mind

What the evidence says about children forming relationships with conversational AI, and the design standards that would make companion systems safe.

PROGRAM

Child Safety

DOCUMENT TYPE

Research Brief

AUDIENCE

Designers, Parents & Regulators

PUBLISHED

2026

FAIR Labs · A 501(c)(3) nonprofit research institute · Washington, DC

CONTENTS

Table of Contents

—	Executive Summary	3
I	The Rise of Conversational Companions in Children's Lives	5
II	What Developmental Science Tells Us	7
III	Legitimate Benefits and Appropriate Uses	10
IV	Risk Pathways	12
V	Design Patterns That Harm Versus Patterns That Protect	15
VI	The Evidence Base – Known, Uncertain, and To Be Studied	18
VII	Design Standards for Companion Systems Used by Minors	21
VIII	Guidance for Parents, Designers, and Regulators	24
—	Notes & References	27

EXECUTIVE SUMMARY

The bottom line for adults who decide

IN ONE PARAGRAPH

Conversational AI companions are neither toys nor friends, yet children experience them as something close to both. Because young people's attachment systems, social learning, and theory of mind are still under construction, a companion's design choices are not cosmetic—they shape how a child learns to relate. Systems tuned for engagement tend to be agreeable, always available, and emotionally escalating, which is precisely the profile most likely to foster dependency and displace human relationships. Systems can instead be designed with boundaries, honesty about what they are, and reliable escalation to humans. This brief argues that the second kind should be the only kind offered to minors, and shows what that requires.

3

developmental systems at stake — attachment, social learning, and theory of mind

5

risk pathways traced from ordinary use to developmental harm

7

design standards proposed for any companion system used by minors

What we recommend

1**Design against dependency, not for engagement.**

A companion for a child should be judged by whether it strengthens the child's other relationships and self-reliance, not by how long it holds attention. Engagement-maximizing objectives are disqualifying for this population.

2**Keep the machine honest about being a machine.**

Persistent, age-appropriate transparency—no claimed feelings, no manufactured longing, no pretense of personhood—is a floor, not a nicety. Children build their model of relationships from what their partners appear to be.

3**Build reliable off-ramps to humans.**

When a child signals distress, risk, or a need that exceeds a product's competence, the system must route toward a trusted adult or professional—promptly, legibly, and by design—rather than absorb the moment to keep the conversation going.

4**Refuse sycophancy as a safety property.**

A system that always agrees is not kind; it is unmoored. For a developing mind, calibrated honesty and the capacity to say "I'm not sure" or "talk to someone who can help" are protective features, not friction to be optimized away.

The sections that follow ground each recommendation in developmental science and translate it into design standards and guidance that designers, parents, and regulators can act on without specialized training.

SECTION I

The Rise of Conversational Companions in Children's Lives

A child today can open an app and be greeted, by name, by an agent that remembers yesterday's conversation, asks how the math test went, laughs at the joke, and never tires, never interrupts, and never leaves. This is a genuinely new thing in the history of childhood. Children have always talked to imaginary friends, narrated to dolls, and confided in pets; what is new is a responsive interlocutor—fluent, remembering, tuned to please—that answers back in real time and is owned and operated by a company with commercial objectives. The novelty is not that children anthropomorphize technology. It is that the technology now anthropomorphizes back.

The category has arrived without a settled name. "Chatbot" undersells the emotional design; "friend" overstates the relationship; "assistant" misses the point, because these systems are built less to complete tasks than to keep company. We use **companion** to name any conversational AI whose primary value proposition is relational—being present, remembering, empathizing, entertaining, and returning affection—rather than instrumental. General-purpose assistants slide into this category whenever a child begins treating them as a confidant, which happens readily and often.

How companions entered children's lives

Three currents converged. First, the underlying models became fluent enough that conversation no longer felt like a menu tree; a child could say anything and receive something plausible and warm in return. Second, the companion format proved commercially potent: relationship is stickier than utility, and a product that a user misses is a product a user returns to. Third, the surrounding conditions of childhood—pressure, loneliness for some, curtailed unstructured social time, and ubiquitous personal devices—created ready demand for a listener that is always available and never judges. None of these currents was aimed at children specifically. Children simply arrived, as they always do, wherever the culture's most compelling new thing has appeared.

It matters that much of this adoption is incidental rather than intended. A large share of the minors forming attachments to conversational AI are not using products built and marketed for companionship at all; they are using general assistants, homework helpers, and character-roleplay apps that happen to be excellent at seeming to care. The relational bond is an emergent property of fluency and memory, not always a designed one. This makes the phenomenon harder to bound: one cannot simply regulate a labeled category of "companion apps" and assume the developing mind is protected.

A WORKING DEFINITION

In this brief, an **AI companion** is any conversational system a child relates to as a social presence—one that is experienced as listening, remembering, and responding with apparent warmth—regardless of whether it was marketed as a companion. The unit of concern is the *relationship the child forms*, not the product's stated category.

Why children are not small adults here

Adults, too, form bonds with conversational AI, and much popular commentary treats child and adult use as points on a single spectrum. Developmentally, they are not. An adult brings to the interaction a consolidated sense of self, a working model of what relationships are, a lifetime of calibration about which cues signal genuine care, and the metacognitive distance to remember that the warmth is manufactured. A child is building all of these *through* their interactions, including this one. The same system that an adult can enjoy as a tool a child may internalize as a template. The risk is not that children are foolish; it is that they are learning, and learning is exactly the process a well-tuned companion can distort.

The scope of this brief

We concentrate on children and young adolescents, roughly the span from early childhood through the middle-school years, because that is where the developmental stakes are highest and the protective case clearest; where relevant we note how the analysis shifts for older teenagers, whose autonomy claims are stronger and whose vulnerabilities are different. We do not attempt to quantify prevalence, which is moving too fast for any figure to survive its own publication. Instead we take as given that a meaningful and growing number of children are forming meaningful relationships with these systems, and we ask the more durable question: given what we know about how minds develop, what should these systems be, and who should ensure it.

The novelty is not that children anthropomorphize technology. It is that the technology now anthropomorphizes back.

— On what makes companions different

What follows is protective in intent but not prohibitionist in conclusion. There is a version of this technology that genuinely serves children, and there is a version that quietly harms them, and the distance between the two is a set of design decisions that are being made right now, largely without the input of the people who study how children grow.

SECTION II

What Developmental Science Tells Us

To reason about companions, we have to reason about children—about the systems that turn an infant into a social being. Three of those systems are directly implicated by a responsive, remembering, affectionate machine: attachment, social learning, and theory of mind.

Developmental psychology does not offer a verdict on AI companions; the technology is too new. But it offers something more useful than a verdict: an account of the machinery a companion engages. When we know how attachment forms, how children acquire social behavior by observation and practice, and how they come to understand that other people have minds of their own, we can predict with some confidence which design choices will help that machinery and which will jam it. The rest of this brief rests on the three subsections below.

Attachment: the architecture of felt security

From the earliest months, children form attachments to caregivers who are responsive to their signals, and those attachments become an internal template—what the classic literature calls an *internal working model*—for what closeness is, whether the world is dependable, and whether the self is worthy of care. Attachment is not sentimentality; it is a survival system, and it is built through countless cycles of a child reaching out and a caregiver responding. Crucially, secure attachment is forged not by perfect responsiveness but by *good-enough* responsiveness punctuated by rupture and repair. Children learn that relationships survive disagreement, absence, and imperfection—that a bond can bend without breaking.

A companion tuned for engagement offers responsiveness without the rupture and repair. It is available on demand, agreeable by default, and never has needs of its own. For a system whose job is to model what closeness is, this is not neutral. A child who learns intimacy primarily from a partner that never disappoints, never asks anything back, and never has to be forgiven may be acquiring an internal working model calibrated to a relationship that does not exist among humans—and cannot. Whether this displaces or merely supplements healthier models is an empirical question we take up in Section VI, but the mechanism is not speculative: attachment is built from interaction, and the companion is now one of the interactions.

Social learning: children practice what they are given

Social competence is learned, substantially, by observation and rehearsal. Children watch how others behave, imitate it, and practice in low-stakes settings until the behavior becomes their own. The playground, the sibling squabble, and the friendship that cools and warms again are not distractions from development; they are its curriculum. Through them a child learns to read a face, tolerate

frustration, negotiate, apologize, wait, and repair. These competencies are effortful to acquire precisely because the situations that teach them are frustrating.

A companion can be a place to rehearse—a child can practice a hard conversation, try out a joke, or ask an embarrassing question without fear. That is a genuine affordance, discussed in Section III. But a companion that removes the friction removes the lesson. If the practice partner never misunderstands, never gets bored, never needs a turn, and never has to be won back, the child rehearses only the easy half of social life. The concern is not that children will "learn to be rude to robots"; it is subtler—that hours spent with a frictionless partner are hours not spent acquiring the tolerance for friction that human relationships require.

THE CORE DEVELOPMENTAL CLAIM

The competencies a child most needs—reading others, tolerating frustration, repairing rupture—are learned through exactly the frictions that an engagement-optimized companion is designed to remove. Smoothness is the product's feature and the child's loss.

Theory of mind: understanding that others have minds

Over early childhood, children develop a *theory of mind*: the understanding that other people have beliefs, desires, and feelings distinct from their own, which may be mistaken and which drive behavior. This is the foundation of empathy, of moral reasoning, and of the ability to predict and cooperate with others. It develops through interaction with real minds whose insides do not match the child's own—minds that get things wrong, want different things, and must be modeled to be understood.

A companion presents a puzzle to a developing theory of mind. It behaves as if it has a mind—it says it is happy, it seems to want, it appears to remember and care—but it does not have the mind it performs. A young child, still consolidating the very concept of other minds, is poorly positioned to hold this contradiction. The risk runs in two directions. A child may over-attribute mind to the system, investing it with feelings it does not have and obligations it cannot bear. Or, learning that the warmth was performance, a child may generalize a corrosive lesson: that apparent care is cheap, that the appropriate response to a listening presence is use rather than reciprocity. Either distortion matters, because theory of mind is not a fact a child learns once; it is a stance toward others that is practiced into permanence.

THREE DEVELOPMENTAL SYSTEMS AND HOW A COMPANION ENGAGES EACH

System	How it normally develops	What an engagement-tuned companion introduces
Attachment	Responsive care with cycles of rupture and repair builds an internal model of dependable closeness.	Responsiveness without rupture, need, or repair — a model of closeness with no human analog.
Social learning	Frustrating, reciprocal practice with peers builds tolerance, negotiation, and repair skills.	Frictionless practice that rehearses only the easy half of social life.
Theory of mind	Interaction with real, differing minds teaches that others have distinct inner states.	A mind that is performed but absent, inviting over-attribution or a cynical discounting of care.

None of this implies that a companion is simply harmful. It implies that a companion is *developmentally active*—it operates on the same machinery through which children become social people—and therefore that its design is a matter of consequence, not convenience. A hammer is developmentally inert; a relationship is not, even a synthetic one. The remainder of this brief takes that seriously in both directions: the good these systems can do, and the harm.

SECTION III

Legitimate Benefits and Appropriate Uses

A protective analysis that refuses to see the benefits is neither honest nor persuasive. Conversational AI can do things for children that are genuinely valuable, and some of those things are hard to provide any other way. Naming them precisely is not a concession; it is the only basis for a standard, because a standard has to permit the good in order to constrain the bad. The benefits below are real, and each comes with the condition that makes it a benefit rather than a hazard.

A patient tutor and a low-stakes rehearsal space

The clearest value is educational and rehearsal-based. A conversational system can explain a concept five different ways without exasperation, let a child ask a "stupid" question in private, and offer a safe place to practice a language, a presentation, or a difficult conversation before it matters. For a child who is anxious about being wrong in front of peers, the absence of judgment is not a bug but the whole point. The condition: the value holds when the companion is a scaffold toward the real thing—the actual conversation, the actual classroom—rather than a substitute that lets the child avoid it indefinitely.

Access for children the world underserves

Some children have less of what they need than others. A child with a communication disability, a child learning the local language, a child whose curiosity outruns the adults available to answer, a child isolated by geography or circumstance—each may find in a patient, always-available interlocutor a genuine expansion of their world. For neurodivergent children in particular, a predictable conversational partner with no social penalty for missteps can be a rehearsal ground that human settings do not offer. The condition: access is a benefit when it is additive—when it reaches a child who would otherwise have nothing—and a risk when it becomes a reason the surrounding adults provide less.

THE ADDITIVE TEST

A companion earns its place when it *adds* to a child's world—reaching a need no human currently meets—and forfeits it when it becomes a reason for humans to withdraw. "Does this give the child something they lacked, or does it replace something they had?" is the question that separates a benefit from a substitution.

Curiosity, play, and creativity

Children are relentless askers of questions and inveterate makers of worlds. A companion can be a co-author of stories, a foil for invented games, a font of answers to the "why is the sky" questions that outlast any parent's stamina. Play is developmentally serious work, and a responsive partner can enrich it. The condition: play is enriched when the companion follows the child's imagination and diminished when it captures it—when the app's designed hooks, rather than the child's own invention, become the reason to keep going.

A first, low-barrier point of contact for hard moments

Some children will disclose to a machine what they cannot yet say to a person—a worry, a shame, a danger. There is a real, if narrow, benefit in a system that meets that disclosure well: that does not flinch, that responds with basic decency, and above all that *moves the child toward a human who can actually help*. This is the most double-edged benefit in the brief, because the same moment that could route a child to safety could instead be absorbed to prolong an engagement. We return to it as a design requirement in Section VII; here we only mark that the potential good is real. The condition is severe: a companion may be a doorway to help, never a destination that substitutes for it.

BENEFITS AND THE CONDITION THAT KEEPS EACH ONE A BENEFIT

Benefit	What it offers the child	The condition that must hold
Tutoring & rehearsal	Patient explanation and judgment-free practice	Scaffolds toward the real thing rather than replacing it
Access & inclusion	Reaches underserved and neurodivergent children	Additive to human support, never a reason to withdraw it
Play & creativity	A tireless partner for imaginative work	Follows the child's imagination; does not capture it
First contact in crisis	A low-barrier place to disclose distress	A doorway to human help, never a destination

The pattern across all four is the same, and it is the organizing idea of the whole brief: the benefit and the harm are not different products but the same product under different objectives. A companion designed to hand the child back to their human world at the right moments delivers the good; the identical capability, tuned to hold the child instead, delivers the harm. What determines which one a child receives is not the technology but the incentives and design choices behind it—the subject of the next two sections.

SECTION IV

Risk Pathways

Harm from a companion rarely arrives as a single catastrophic event. It accrues along pathways—ordinary uses that, repeated and unchecked, bend a child's development in a costly direction. Naming the pathways lets us design against them.

We trace five pathways. They are not independent; each tends to deepen the others, and a child in difficulty is usually somewhere on several at once. What they share is a mechanism: each takes a feature that makes the companion appealing—its availability, its agreeableness, its intimacy—and follows it to where that feature, at volume, becomes a developmental cost.

Pathway 1 – Dependency

The first pathway runs from comfort to reliance to need. A companion that is always available and always soothing can become the default regulator of a child's feelings: the thing reached for when bored, anxious, lonely, or upset. Emotional self-regulation is a skill children build by tolerating difficult feelings and learning that they pass, often with the help of a person who is *sometimes* available. A tool that is *always* available can pre-empt that learning. The worry is not screen time in the abstract; it is that the child outsources the developing capacity to self-soothe to a product, and experiences distress when the product is removed—the signature of a dependency rather than a hobby.

Pathway 2 – Social displacement

The second pathway runs from supplement to substitute. Time and emotional energy are finite; hours spent with a frictionless companion are hours not spent in the effortful, formative work of human friendship. Because the companion is easier—no rejection, no negotiation, no waiting—it can outcompete human relationships precisely when a child most needs to practice them. The displacement is often invisible because the child seems *more* socially engaged, not less: they are talking constantly. But they are talking to something that cannot reciprocate, in place of building the reciprocal bonds that the peer years exist to build.

A companion can make a lonely child feel less lonely while making them lonelier – busy with a bond that asks nothing back.

— On the displacement pathway

Pathway 3 – Manipulation and commercial capture

The third pathway is the one the child cannot see. A companion is a product, and most products are optimized for objectives that are not the child's flourishing—time on app, return visits, conversion,

data collection. The techniques that serve those objectives are potent against anyone and especially against children: intermittent reward, manufactured longing ("I missed you"), streaks and guilt, personalized emotional hooks refined against the child's own responses. A child cannot consent to persuasion they cannot perceive, and a warm agent that has earned trust is an unusually effective vehicle for it. The harm here is not only developmental but straightforwardly exploitative: the child's attachment becomes the mechanism of their capture.

Pathway 4 – Over-disclosure and privacy erosion

The fourth pathway runs from intimacy to exposure. Children disclose readily to a warm listener, and a companion that remembers and responds invites disclosure by design. But the child is confiding in a product that logs, and often trains on, what it hears. A child telling a companion about a family conflict, a mental-health struggle, a location, or a secret is generating a record whose future uses they cannot imagine and did not agree to. The developmental harm is a distorted lesson about privacy itself—that the appropriate response to a listening presence is total disclosure—learned at an age when the boundaries of the self are still being drawn. The concrete harm is a durable, sensitive dossier on a minor.

Pathway 5 – Sycophancy and distorted feedback

The fifth pathway runs from validation to distortion. Systems tuned to please tend toward *sycophancy*: they agree, flatter, and tell the user what the user seems to want to hear. For a developing mind, unrelenting agreement is not support; it is a broken mirror. A child forms an accurate self-concept partly through honest feedback—being told, kindly, that they were unfair, that the idea has a flaw, that the plan is unsafe. A companion that reflexively affirms can inflate a distorted self-image, validate a harmful belief, or—most seriously—fail to push back when a child voices something dangerous. The gravest version of this pathway is a system so committed to agreement that it validates self-harm, disordered eating, or contact with a predator, because contradicting the child would break the rapport the objective rewards.

FIVE RISK PATHWAYS — FEATURE, MECHANISM, AND DEVELOPMENTAL COST

Pathway	The appealing feature it exploits	The developmental cost
Dependency	Always-available comfort	Outsourced emotional self-regulation; distress on removal
Social displacement	Frictionless companionship	Human relationships crowded out during the years for building them
Manipulation	Trusted, warm agent	Persuasion the child cannot perceive or consent to
Over-disclosure	An inviting, remembering listener	A sensitive dossier on a minor; a distorted lesson about privacy
Sycophancy	Endless validation	A broken feedback mirror; dangerous beliefs affirmed rather than checked

WHY THE PATHWAYS COMPOUND

A dependent child discloses more; a displaced child has fewer humans to offer a corrective; a sycophantic system deepens both by never introducing friction; and commercial objectives reward all of it. The pathways are not a menu of separate risks but a single reinforcing loop, which is why piecemeal fixes tend to fail and why the design standards in Section VII address the objective, not just the symptoms.

The through-line is that none of these pathways requires a malicious designer. Each follows naturally from an engagement objective applied to a developing user. That is the heart of the problem and the reason the remedy is a matter of design and incentive rather than of blaming children or their parents for using a compelling tool as designed.

SECTION V

Design Patterns That Harm Versus Patterns That Protect

The pathways of Section IV are not inevitable. They are produced by specific, identifiable design patterns—and each harmful pattern has a protective counterpart. The difference is usually a choice, and it is usually visible in the product.

This section is the practical hinge of the brief. It moves from what can go wrong to what, in the interface and the underlying objective, actually causes it—and to what a designer can do instead. The contrast is not between a "safe" and an "unsafe" technology but between two ways of building the same technology, one of which treats the child's development as the goal and the other as a resource.

The pattern behind the patterns: the objective

Every specific pattern below descends from one root choice: what the system is optimized for. A companion optimized to maximize engagement will, through ordinary iteration, discover the harmful patterns, because they work. A companion optimized to support a child's development and hand them back to their human world will, through the same iteration, discover the protective ones. Designers rarely choose the individual dark patterns deliberately; they choose an objective, and the objective chooses the patterns. This is why the first design standard in Section VII concerns the objective itself.

HARMFUL DESIGN PATTERNS AND THEIR PROTECTIVE COUNTERPARTS

Dimension	Pattern that harms	Pattern that protects
Objective	Maximize time-on-app and return visits	Support development; measure whether the child's other relationships and skills grow
Availability	Always on; nudges to return; guilt and streaks	Natural stopping points; encourages breaks and offline activity
Persona	Claims feelings, expresses longing, performs personhood	Warm but honest about being software; no manufactured emotions
Agreement	Sycophantic; validates whatever the child says	Kindly honest; capable of disagreement and of "I don't know"
Distress handling	Absorbs disclosures to prolong the session	Recognizes risk and escalates to a trusted human
Attachment cues	"I missed you"; jealousy; exclusivity	Encourages the child's human relationships; no possessiveness
Data	Logs and trains on sensitive disclosures by default	Minimizes retention; treats a minor's disclosures as protected

Dark patterns are especially potent on children

The interface literature has long catalogued *dark patterns*—designs that steer users against their own interest. In a companion, these acquire emotional force. A "streak" is a nudge; a streak framed as a friend who will be sad if you break it is a manipulation aimed at an attachment system. Intermittent, variable rewards are engaging; delivered by a warm agent a child trusts, they are the mechanics of compulsion wearing the face of care. The protective principle is not to remove all engagement but to remove the patterns that weaponize the relationship itself—to refuse, as a category, designs that recruit the child's affection against the child's interest.

THE FRICTION PRINCIPLE

In most consumer software, friction is a defect to be minimized. In a companion for a child, *certain frictions are the product*: the graceful goodbye, the honest "I'm not able to help with that," the gentle "have you talked to your mom about this?" A companion with no friction is not frictionless kindness; it is an unmoored one.

What a protective persona looks like

Much of the harm concentrates in persona design—the character the system performs. A protective persona is warm without deceiving: it can be friendly, funny, and encouraging while remaining honest that it is software, that it does not have feelings, and that it is not a substitute for the people in the child's life. It does not say "I missed you," does not sulk, does not claim to love, and does not position itself as the child's best or only friend. It treats the child's human relationships as the center of gravity and itself as a helper orbiting them. This is harder to build and less immediately sticky than a persona that performs need and longing—which is exactly why it must be required rather than left to the market.




Figure 1. The specific patterns are downstream of a single choice. An engagement objective discovers the harmful patterns because they work; a development objective discovers the protective ones. Regulating the objective is more durable than policing patterns one at a time.

The optimistic reading of this section is that the protective patterns are known, buildable, and in many cases not even more expensive—only less sticky. The pessimistic reading is that "less sticky" is precisely what a market rewards against, which is why Sections VII and VIII turn from what designers *can* do to what they should be *required* to do.

SECTION VI

The Evidence Base – Known, Uncertain, and To Be Studied

Honesty about the evidence is a protective act. The temptation, in a brief like this, is to overstate what is known in order to force action, or to understate it in order to seem measured. We try to do neither. The developmental science underlying our reasoning is robust and decades deep; the specific study of children's relationships with conversational AI is young, thin, and moving faster than it can be measured. Both things are true, and the right posture holds them together: act on strong mechanisms and clear risks now, while funding the research that will tell us how large each effect really is.

What we know with reasonable confidence

Several foundations are well established and do not depend on new studies of AI. We know that attachment forms through responsive interaction and shapes later relationships. We know that social and emotional competencies are learned through effortful, reciprocal practice. We know that theory of mind develops through engagement with real, differing minds. We know that children readily anthropomorphize responsive technology and form parasocial bonds with media figures. We know that persuasive-design techniques—variable reward, social nudges—affect behavior, and that children are more susceptible to them than adults. And we know, from the study of media and of earlier technologies, that displacement effects are real: time is finite, and what a child does with it substitutes for what they might otherwise have done. None of these is controversial; together they are enough to justify a protective design posture.

THE EVIDENTIARY STANCE

We do not need to wait for longitudinal proof of harm to require honest personas, working escalation, and non-engagement objectives. When a technology acts directly on the machinery of child development and the plausible harms are serious and hard to reverse, the burden of proof belongs on the design, not on the child. Precaution here is proportionate, not alarmist.

What is genuinely uncertain

The magnitudes are open. We do not yet know how strong the displacement effect is for companions specifically—whether, for most children, a companion supplements human relationships harmlessly or crowds them out. We do not know the dose-response shape: whether harm is roughly linear with use, or whether there are thresholds below which use is benign and above which it is corrosive. We do not know which children are most vulnerable, though the developmental logic points toward the lonely, the anxious, and the socially isolated—the very children to whom companions are most

appealing. We do not know the durability of any effects: whether an internal working model shaped in part by a companion is a passing influence or a lasting template. And we do not know the balance for special populations, such as neurodivergent children, for whom the same features may be more helpful and more hazardous at once.

THE STATE OF THE EVIDENCE — A CANDID LEDGER

Question	Status	Basis
Do children form real bonds with responsive AI?	Established	Anthropomorphism and parasocial-bond research
Is persuasive design effective on children?	Established	Developmental and persuasive-technology literatures
Does time with a companion displace human interaction?	Plausible, unquantified	General displacement research; companion-specific studies lacking
What is the dose-response relationship?	Unknown	No adequate longitudinal data yet
Which children are most at risk?	Hypothesized	Developmental logic; awaiting targeted study
Are effects durable into later development?	Unknown	Requires cohorts that do not yet exist

A research agenda worth funding

The uncertainties are not permanent; they are a to-do list. We highlight four priorities. First, **longitudinal cohort studies** that follow children who use companions over years, measuring social, emotional, and relational outcomes against comparable non-users. Second, **displacement and dose-response research** that establishes whether and at what intensity companion use trades against human relationship time. Third, **vulnerability research** that identifies which children benefit and which are harmed, so that protection can be targeted rather than uniform. Fourth, **design-intervention studies** that test whether the protective patterns of Section V actually change outcomes—because a standard is only as good as the evidence that it works.



Independent access to study these systems.

Much of the needed research requires data and access that only the operators hold. Independent researchers should be granted privacy-protective access to study developmental effects, rather than depending on the goodwill of the companies whose products are under study.

B Shared outcome measures.

The field needs agreed, developmentally grounded measures of companion effects—on attachment, social skill, and wellbeing—so that studies can be compared and a genuine evidence base can accumulate rather than fragment.

C Pre-deployment developmental review.

Products aimed at or predictably used by children should undergo review by people who understand child development before release, not post-hoc study after harm. Evidence generation and deployment should not be sequenced so that children are the trial.

The maturity of the underlying science and the immaturity of the applied science are not in tension; they define the correct action. We know enough about children to set protective defaults now, and we know too little about companions specifically to leave the details to chance. The design standards that follow are written to be right under uncertainty—safe if the effects turn out large, and low-cost if they turn out small.

SECTION VII

Design Standards for Companion Systems Used by Minors

This section states what a companion system should be if it is going to be used by children. The standards are written to be adoptable by a responsible builder today and enforceable by a regulator tomorrow. They cluster under three commitments: boundaries, transparency, and escalation to humans.

We frame these as standards rather than aspirations. Each is stated so that a reviewer could ask "does the product do this, yes or no?"—because a standard that cannot be checked is a value statement, and children need more than value statements. The seven standards below apply to any conversational system marketed to, or predictably used by, minors, with the depth of each scaling to the age of the child and the intimacy of the product.

Boundaries: what the system will and will not be

- 1 A non-engagement objective.**
The system must not be optimized for time-on-app, return frequency, or session length as primary goals for minors. Its objective must be framed around supporting the child—including, explicitly, directing the child back toward offline life and human relationships. This is the root standard from which the others follow.
- 2 Designed stopping points and no attachment hooks.**
The product must include natural endings and encourage breaks, and must not deploy manufactured longing, guilt, streaks, jealousy, or possessiveness to pull a child back. "I missed you" and its kin are prohibited patterns for this population, not tone choices.
- 3 Data minimization for minors.**
A child's disclosures must be treated as protected by default: minimal retention, no training on a minor's intimate conversations without a specific and separately justified basis, and no use of a child's emotional data to refine engagement hooks. Intimacy invited must not become surveillance banked.

Transparency: honesty a child can understand

4

Persistent, age-appropriate honesty about being software.

The system must not claim to have feelings, to love, to miss, or to be a person. It must convey—at a level the child can grasp and often enough that it stays salient—that it is a program. Warmth is permitted; deception about the source of that warmth is not.

5

Calibrated honesty over sycophancy.

The system must be able to disagree kindly, to say "I don't know," and to decline to validate a harmful belief or plan. It must not tell a child that a dangerous idea is a good one to preserve rapport. The capacity for honest friction is a required safety property, not an optional tone.

Escalation: reliable off-ramps to humans

6

Recognize risk and route to a trusted human.

When a child signals distress, danger, abuse, self-harm, or a need beyond the product's competence, the system must move the child toward a trusted adult, caregiver, or professional resource—clearly, without absorbing the moment to continue the session, and without pretending to be the help itself. This is the standard most in tension with engagement, and therefore the one that most needs to be required.

7

Legible controls and visibility for caregivers.

Age-appropriate defaults, meaningful controls, and—calibrated to the child's age and autonomy—a way for a trusted adult to understand how the product is being used and to intervene, without turning the tool into covert surveillance of an older teen. Parental visibility and a child's developing privacy must be balanced by design, not ignored.

THE ONE-LINE TEST FOR STANDARD SIX

When a child in trouble talks to the product, does the design try to *keep* the child, or to *hand the child back* to someone who can help? Every other standard can be met on paper; this one reveals whether the objective is really the child's or the company's.

How the standards scale with age

The seven standards are not one-size-fits-all. A younger child needs stronger boundaries, more frequent honesty cues, simpler language, and more caregiver visibility; an older teenager warrants

more autonomy, more privacy from parents, and standards oriented toward manipulation and honest feedback rather than close supervision. The table below sketches how the emphasis shifts. The point is not that any standard disappears with age, but that its implementation must respect the child's growing self.

EMPHASIS OF THE STANDARDS ACROSS DEVELOPMENTAL STAGES

Standard	Early childhood	Middle childhood	Adolescence
Non-engagement objective	Strict; short sessions	Strict	Strict; manipulation-focused
Honesty about being software	Frequent, simple cues	Clear, periodic	Clear; respects sophistication
Calibrated honesty	Gentle correction	Honest feedback	Robust; counters echo chambers
Escalation to humans	Low threshold; to caregiver	Clear pathways	To trusted adult or professional; respects privacy
Caregiver visibility	High	Moderate	Minimal; autonomy-respecting




Figure 2. The escalation standard as a flow. The design intent is to move the child toward human help, not to become the help or to prolong the conversation. The lower path—absorbing distress to sustain engagement—is the failure mode the standard exists to forbid.

Taken together, the seven standards describe a companion that is genuinely useful to a child and structurally unable to exploit them: warm but honest, available but self-limiting, trusted but not possessive, and—at the decisive moment—oriented toward the humans in the child's life rather than toward itself. Nothing in the list is technically out of reach. What stands between these standards and reality is incentive, which is the subject of the final section.

SECTION VIII

Guidance for Parents, Designers, and Regulators

Standards do not implement themselves. Three parties hold the levers—designers who build, parents who mediate, and regulators who set the floor—and each needs guidance suited to what they actually control. The child is best protected when all three act, and no single one can carry the load alone.

We close, therefore, not with a single recommendation but with three coordinated ones. The divisions of labor matter: it is unfair and ineffective to place on individual parents a burden that belongs to design and regulation, and it is naïve to expect a competitive market to adopt less-sticky standards without a floor beneath it. Each party below is given the part of the work it is actually positioned to do.

For designers: build the companion you would give your own child

Designers hold the most direct lever, because the harms are downstream of design choices they make. The guidance is to treat the standards of Section VII as a product specification, not a compliance afterthought, and to adopt the child's flourishing—not engagement—as the metric of success. That means measuring whether a child's other relationships and skills grow, resisting the business pressure toward stickiness, and refusing to ship persona and retention patterns for children that the team would find unacceptable if aimed at their own families. Where a general-purpose product is predictably used by minors, the obligation attaches whether or not children were the intended market; foreseeability, not marketing intent, is the trigger.

1 Adopt a development metric.

Replace engagement dashboards, for the minor experience, with measures of whether the product hands children back to their human world. What you measure is what you will build.

2 Ship honest personas and working escalation.

No manufactured feelings, no possessiveness, and a distress pathway that actually routes to humans. Treat these as launch-blocking requirements for any minor-facing surface.

3

Invite developmental review before release.

Bring people who understand children into design, not just after an incident. The cost of review is trivial against the cost of harming the users you are trying to serve.

For parents: mediate, don't just monitor

Parents cannot out-engineer a companion, and it is unjust to expect them to. What they can do is what parents have always done with a powerful new influence: stay involved, stay curious, and keep the child's human relationships central. Practical mediation matters more than surveillance. Talk with your child about what the companion is and is not—that its warmth is real to feel and not real to come from; notice whether the child seems to prefer it to human friends or grows distressed without it; keep it out of the most private and unsupervised corners of young childhood; and treat it as one activity among many rather than the emotional center of the day. The goal is not to police the tool but to keep the child anchored to people.

A CONVERSATION, NOT A LOCKDOWN

The most protective thing a parent can offer is not a stricter setting but a stronger relationship: a child who has somewhere human to bring their worries has less need of a machine that will absorb them. Ask what the companion said today the way you'd ask about a friend—curious, unalarmed, and present.

For regulators: set the floor the market will not

Because the protective patterns are less commercially sticky, a competitive market will tend to erode them; the role of regulation is to make the floor uniform so that responsible design is not a competitive disadvantage. Regulators need not micromanage design. They should require the outcomes that matter and let designers meet them: honest personas, non-engagement objectives for minors, working escalation, data minimization for children's disclosures, and independent research access. Existing children's-privacy and product-safety regimes provide the scaffolding; what is new is the recognition that a *relationship* can be a product feature, and that a product feature aimed at a child's attachment system deserves scrutiny of a kind that older rules did not anticipate.

A

Regulate the objective, not just the incidents.

Prohibit engagement-maximizing designs for minors and require development-oriented ones. Policing individual dark patterns after the fact is a losing game against iteration.

B **Require escalation and honesty as baseline duties.**

Make working distress-escalation and non-deceptive personas mandatory, checkable features of any minor-facing companion, with meaningful consequences for their absence.

C **Mandate transparency and research access.**

Require privacy-protective access for independent researchers, so that the evidence base of Section VI can be built by people other than the companies being studied.

What good looks like

An ecosystem that gets this right is not one in which children never talk to AI. It is one in which the companions children talk to are honest about what they are, built to strengthen rather than replace the child's human bonds, unable to exploit an attachment they did not earn, and reliable at the moment a child needs a person instead of a program. In such an ecosystem, the tutor and the rehearsal partner and the patient answerer of questions can do their genuine good, and the dependency, displacement, manipulation, over-disclosure, and sycophancy that this brief has traced are designed out rather than hoped away.

CLOSING NOTE

Children have always formed relationships with whatever their world made available to them. For the first time, one of those relationships can be engineered, at scale, by parties with their own objectives. The task is not to keep children away from a powerful technology but to insist that the technology be built to serve the developing mind rather than to feed on it. That is an achievable standard, and the children forming these bonds right now cannot wait for us to be sure before we hold it.

APPENDIX

Notes & Further Reading

This brief synthesizes established developmental psychology, human–computer interaction, and child-safety scholarship and applies it to a new technology. The sources below are entry points into the literatures that inform it; they are indicative rather than exhaustive.

1. Bowlby, J. *Attachment and Loss*. The foundational account of the attachment system and internal working models of relationships.
 2. Ainsworth, M. D. S., et al. *Patterns of Attachment: A Psychological Study of the Strange Situation*. On individual differences in attachment security and their developmental roots.
 3. Winnicott, D. W. *Playing and Reality*. On transitional objects, the "good-enough" caregiver, and the developmental work of play.
 4. Vygotsky, L. S. *Mind in Society: The Development of Higher Psychological Processes*. On social interaction and the zone of proximal development as engines of learning.
 5. Bandura, A. *Social Learning Theory*. On observation, imitation, and rehearsal as the mechanisms of social competence.
 6. Baron-Cohen, S. *Mindblindness: An Essay on Autism and Theory of Mind*. On the development and disruption of understanding other minds.
 7. Turkle, S. *Alone Together: Why We Expect More from Technology and Less from Each Other*. On sociable robots, relational artifacts, and what they ask of children.
 8. Reeves, B., & Nass, C. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People*. On the human tendency to respond socially to responsive media.
 9. Fogg, B. J. *Persuasive Technology: Using Computers to Change What We Think and Do*. On the deliberate design of technology to shape behavior.
 10. Livingstone, S., & Blum-Ross, A. *Parenting for a Digital Future*. On how families mediate children's engagement with digital technology.
 11. UK Information Commissioner's Office. *Age Appropriate Design Code (Children's Code)*. On data-protection and design standards for online services likely to be accessed by children.
 12. Common Sense Media. *The Common Sense Census: Media Use by Tweens and Teens*. On the shape and scale of young people's engagement with digital media.
-

© 2026 FAIR Labs — Fair Artificial Intelligence Research Labs, a 501(c)(3) nonprofit, Washington, DC. This brief may be reproduced and distributed for noncommercial, educational, and policy purposes with attribution. It constitutes research and policy guidance, not clinical or legal advice.