The Googly-Eye Principle for Robots in Autism Therapy

In autism therapy, progress often comes from repeated opportunities to practice social interaction. By offering clear, low-stakes invitations to engage, even simple robots can scaffold social learning in everyday life.

By Rebecca Ramnauth

DOI: 10.1145/3812682

In our house, maintaining a ready supply of googly eyes was not unusual. My parents stuck them on everything from the lemons on our kitchen counter to the caps we wore at baseball practice. We woke to an alarm clock that could "see" our reluctance on early school mornings and at night we thanked our toothbrush holder for its patience as we flossed. What appeared to be a playful family tradition was, in fact, a deliberate intervention: an effort to embed frequent, naturalistic opportunities for my sister to practice making eye contact.

My sister was diagnosed with autism two decades ago, at a time when diagnostic frameworks and available supports differed substantially from those of today. Autism Spectrum Disorder (ASD) is a neurodevelopmental condition characterized by differences in social communication and interaction, often accompanied by restricted interests or repetitive patterns of behavior.

The criteria used to define autism have evolved over time—sometimes reflecting shifts in clinical perspective and service infrastructure as much as new mechanistic understanding. As a spectrum, autism manifests in diverse ways. A person may possess a skill ("I can do it."), struggle to deploy it reliably in certain contexts ("I won't do it now."), or actively avoid it because it is uncomfortable or overwhelming ("I don't like to do it."). This variability is reflected in how individuals communicate and engage socially. Some individuals are highly verbal; others communicate minimally or through alternative modalities. Some actively seek social connection but find it difficult to interpret subtle cues; others experience social participation as cognitively taxing or affectively intense. Despite this heterogeneity, many autism interventions rely on a remarkably consistent principle: Provide frequent, naturalistic opportunities to practice specific social skills.

Building on this foundation, pedagogical approaches vary along several surface dimensions. Instruction may be explicit, involving verbal modeling or corrective feedback, or more implicit. Discrete skills may be targeted in isolation—such as initiating eye contact or responding to one's name—or embedded within broader, contextually integrated competencies such as sustaining conversation or regulating emotion. Delivery may be directed by parents or therapists or mediated by peers. Still, across these traditions, the underlying logic is consistent: Increasing the density of low-stakes opportunities for skill expression makes desired behaviors more durable and more likely to transfer across contexts.

In retrospect, my parents' "googly-eye" intervention embodied this same

Image by Anton Vierietin / Shutterstock.com

logic and, in many ways, anticipated my later work in robotics (see Figures 1 and 2). Robots can provide on-demand, personalized interactions that are both socially situated and physically co-present. Unlike disembodied or purely virtual technologies, they can turn to face you, pause for your response, and hold shared space in ways that invite social response and reinforce normative patterns of interaction. In this sense, robots offer a distinct advantage—a low-stakes yet socially evocative partner through which social skills can be practiced within everyday contexts.

One of the earliest long-term, in-home robotic interventions for autism used a minimally expressive robot to model appropriate gaze behavior [1]. In comparison, our robot would turn toward the child to invite eye contact, then shift its attention to the parent, encouraging the child to follow its gaze [2]. Although the behavior itself was simple, the interaction engaged several foundational components of social development targeted in autism therapy (making eye contact, following another's gaze, and ultimately sharing attention). By delivering thousands of these small, structured invitations, the robot increased the density of opportunities for children to practice these skills within everyday family interactions at home (see Figure 1). Within the first two weeks of our robot study, children with autism showed sustained improvements in spontaneously initiated eye contact and joint attention [2].

PRINCIPLED DESIGN IN THE ABSENCE OF MECHANISTIC CLOSURE

Over the past two decades, hundreds of studies have examined the use of robots in autism interventions, reporting increased engagement, improved attention regulation, and gains in social behaviors such as conversational reciprocity and spontaneous imitation when robots are incorporated into interaction [3, 4]. Although these positive effects are consistently reported across diverse contexts and robot designs, the field lacks a consolidated account of why robots are effective in eliciting engagement and observable expressions of social skill in autism.

Several compelling hypotheses have been proposed. While each has intuitive appeal, each also invites counterargument. One popular account situates robots along a "spectrum of sociability." On one end are inanimate toys that elicit little or no social behavior whatsoever; on the other are humans, whose social cues are rich, nuanced, and at times overwhelming. Robots, in this view, occupy a middle ground—social enough to invite interaction, yet simplified enough to remain cognitively tractable. Although appealing, this framing rests on two assumptions: social complexity varies along a linear continuum, and humans represent its maximal point—assumptions that are neither theoretically nor empirically secure.

Another explanation emphasizes legibility. Robotic cues (such as their gaze shifts, gestures, and vocal prosody) are explicitly contingent, sometimes exaggerated, and often temporally discrete. Compared to the complex and overlapping signals exchanged between humans, these cues are presumed to be easier to interpret and therefore easier to respond to. Yet this account overlooks an important complication. For many users—especially young children, the population most extensively studied in autism research—robots are novel agents whose status is ambiguous. This ambiguity can arise at multiple levels: individual features ("Is the moving circle on its screen an eye or a nose?"), ontological category ("Is it a tool or a friend?"), the normative rules that follow ("It shouldn't do that."), and even meta-reasoning about its legitimacy ("Does this thing even get to decide?"). Robotic signals may not be inherently clearer; they are simply unfamiliar. Interpreting them therefore requires a layered chain of inference that is not routinely exercised in everyday human interaction.

Social skills develop through repeated, low-stakes opportunities for engagement embedded within everyday life.

Perhaps, instead, robots offer a form of "social neutrality," wherein their absence of social status or prior relational history makes them less intimidating and more approachable than human peers. To this, one may argue that robots are not culturally neutral artifacts. Public narratives surrounding AI and robotics—ranging from privacy concerns to dystopian imaginaries—shape how these systems are perceived long before interaction begins. From a roboticist's perspective, these narratives influence how agency is framed, how autonomy is constrained, how data practices are made visible, and how ongoing consent is structured within the interaction. From a researcher's perspective, they shape recruitment itself—who feels comfortable enrolling in a study, who opts out, and whose experiences ultimately inform the literature.

Despite the breadth of experimentation on robots for autism therapy, explaining why these systems are effective remains methodologically challenging. Autism is highly heterogeneous, and clean control conditions are difficult to construct. Comparative human-versus-robot designs often introduce relational differences that are difficult to disentangle, and robot-absent baselines are not always feasible within socially contingent interactions. As a result, a comprehensive mechanistic account remains elusive.

Still, the absence of a definitive mechanism does not preclude principled design. The field continues to generate hypotheses because each offers a plausible design rationale—what to simplify, what to emphasize, what cues to foreground, and how to structure engagement. These operational assumptions, in turn, influence both how systems are built and how their effectiveness is evaluated. While we should be careful not to mistake recurring design heuristics for explanatory theory, progress in building robots that support autism therapy does not depend on achieving full mechanistic closure.

Figure 1. A minimally expressive robot scaffolds joint attention in an in-home autism intervention.

The system was deployed in family homes for daily 30-minute sessions over a one-month period with children with autism (ages 6–12). Across the deployment, 269,278 gaze events were captured and analyzed, revealing significant increases in spontaneous eye contact and joint attention between children and their caregivers as the robot repeatedly modeled appropriate gaze shifts [2].

ADDITIONAL DESIGN COMMITMENTS

Even as the field continues to investigate these mechanisms, we can ask a complementary question: Given what we observe in practice, how should robots be designed to support autism therapy responsibly and effectively? In other words, if the mechanisms remain unsettled, let us articulate design principles that remain robust across competing explanations. In what follows, I outline three design commitments that have guided my work in robot-assisted autism therapy. Reminiscent of the hypotheses previously discussed, these principles prioritize: 1) decentering human likeness, 2) designing for interpretive tractability, and 3) embedding interaction within real-world social contexts.

These commitments are not intended as universal or definitive solutions. Instead, they reflect choices that have emerged through building and deploying real systems and through what we currently understand about autism. They are offered as principled guides for practice rather than as a comprehensive framework.

Decoupling sociality from human likeness. There is something intuitively Frankenstein-ian about building a robot, as if constructing a novel social partner were akin to constructing a human being. Indeed, an active line of research pursues robots that are highly anthropomorphic with silicone-based skin and expressive facial musculature, limbic gestures that are human-like and coordinated, complete with secondary features such as eyebrows and lower eyelids. More broadly, across AI and robotics, human likeness often functions as an implicit benchmark. Early natural language systems were evaluated by whether a computer could be mistaken for a human in the Turing Test, and contemporary systems are frequently assessed relative to human performance baselines.

Figure 2. Robots supporting social and emotional skills across diverse real-world contexts.

My work extends the same design philosophy across multiple domains of social learning and regulation. Shown here are representative studies in which robots reduce isolation among children during pandemic lockdowns [8], support deep-breathing and anxiety regulation through haptic feedback, support job-relevant skills such as resilience to workplace interruptions [9] and conversational competence for adults with autism across a wide range of functional abilities [10], and assist with emotional de-escalation in school settings [11]. Across these varied applications, each system provides structured, low-stakes opportunities to practice valuable social skills within everyday environments.

However, success need not be indexed to human similarity. If we set aside the human as the primary referent, we can envision robots that are effective not by approximating human behavior, but by occupying interactional roles that humans do not easily fill. Even without comparative experimental trials, it is not difficult to imagine why an animatronic teddy intuitively invites calming, tactile interactions in ways a human peer cannot. Further, a small, limbless robot can be less intimidating than a life-sized humanoid that stands face-to-face in height. A minimalist interface that exaggerates gaze cues may elicit clearer attention-following than the nuanced signals of a human face. Likewise, a stateless machine can offer a form of consistent, non-judgmental presence that human partners—whose interactions accumulate memory and social history—cannot sustain indefinitely.

There is a subtle parallel to this observation in autism interventions, where neurotypical behavior is often treated as the normative reference point for successful interaction. Yet the goal of autism interventions is not to replicate neurotypical interaction in full fidelity. Rather, it is to support learning and engagement within environments that can be calibrated to the needs of the learner. From this perspective, the most effective robot may not be the one that most faithfully simulates a human partner, but the one that selectively restructures aspects of social interaction in ways that humans cannot consistently maintain.

Seen in this light, robots are better understood not as artificial humans, but as interactional scaffolds. Their value lies in how they structure opportunities for engagement: cues can be exaggerated without awkwardness, behaviors repeated without fatigue, and social responses delivered with a consistency that human partners rarely sustain. The relevant design question therefore shifts. Instead of asking how closely a robot should resemble a human, we can ask what interactional affordances it should provide to support learning and participation in the broader social world.

Treating minimalism as affordance. The reason we build robots is to help people. In practice, a robot's helpfulness is expressed through descriptors such as being useful, productive, supportive, or even entertaining. Under this broad notion, research has produced a wide range of helpful systems: machines built for functional task assistance, socially interactive partners offering developmentally attuned support, and passive devices intended for on-demand entertainment. Across these efforts, we nonetheless evaluate system success as external to the system itself. The metrics most commonly used are perceptual (what users think), behavioral (how users act in response to the robot), relational (whether engagement is established and sustained), and outcome-based (as judged by people). In this sense, a robot's value does not reside solely in its internal capabilities, but in how its behavior is understood and situated within human activity.

If the usefulness of a robot is interpreted, then it does not necessarily scale with technical sophistication. Increasing behavioral repertoire and intelligence does not automatically produce a more effective robot. In many contexts, surprisingly little is required for a robot to be useful and effective.

A number of existing systems illustrates this principle. A single-eye robot that alternates its gaze between shared task content and human partners has been shown to produce clinically verified gains among children with autism [2]. A spherical robot that simulates breathing through a single degree of freedom has been demonstrated to reduce anxiety in pediatric waiting rooms [5]. Outside explicitly therapeutic settings, household vacuum robots with no explicit social affordances nonetheless inspire attachment and continued value as companions even when defunct [6]. Likewise, a small stationary robot with no limbs has proven sufficient to motivate sustained daily exercise in longitudinal deployments [7]. Across these examples, the systems themselves are apparently modest, yet they succeed because their behaviors are easily interpretable and readily embedded within everyday routines.

Nevertheless, the broader trajectory of robotics and AI increasingly emphasizes the development of general-purpose systems. Recent advances in large foundation models have accelerated efforts to build robots that can integrate perception, language, and action into unified architectures capable of operating across diverse tasks and environments. On one hand, these developments promise unprecedented breadth of capability. On the other hand, they expand the space of interpretations users may reasonably make about the system's abilities.

For example, a robot that responds fluently to personal disclosures may be assumed to understand emotional nuance. Similarly, a system that offers advice across a wide range of situations may be treated as a reliable authority, then trusted to handle critical decisions. As robots exhibit broader and more integrated behaviors, they can invite expectations that the underlying system cannot reliably sustain. When those expectations are violated, interaction can quickly break down—not necessarily because the robot has failed technically, but because the user's interpretation has outpaced what the system can actually support.

Figure 3. Real-world perception during an in-home deployment.

The robot estimates head pose and facial landmarks to infer user attention and engagement in a cluttered home environment [2]. When deployed in real-world settings for naturalistic interaction alongside multiple people, robots require far greater technical robustness than systems operating in controlled laboratory settings.

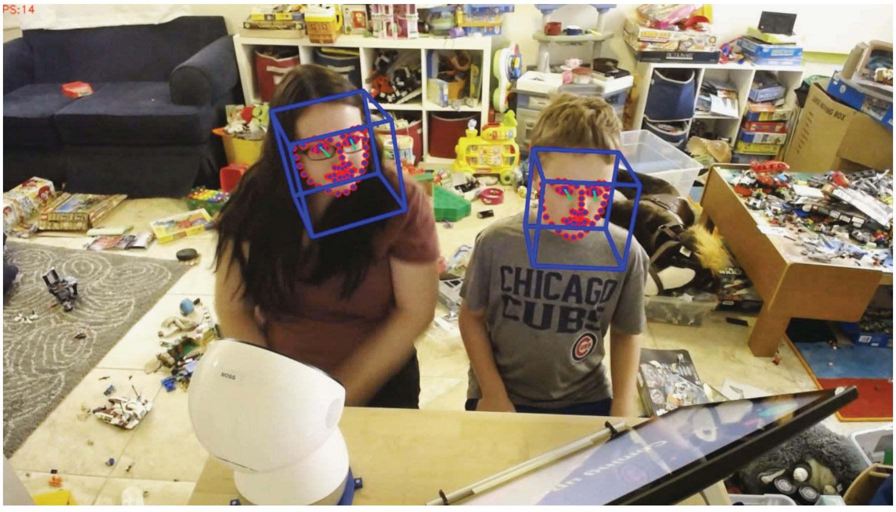

One way to mitigate this problem is through deliberate constraint. Systems with limited behavioral repertoires can communicate clearer boundaries about what they can and cannot do. Their reduced expressivity narrows the range of plausible interpretations, helping stabilize expectations and reduce the likelihood of interpretive failure. These considerations are particularly relevant in autism interventions, where reduced complexity of social cues can be essential for achieving therapeutic goals. Minimalism therefore functions not as a limitation, but as an affordance.

Creating invitations to the real world. Minimalist interaction design should not be mistaken for technical simplicity. Even robots that present only a small set of behaviors must still operate reliably within complex, unpredictable human environments. In therapeutic contexts especially, the effectiveness of a robot does not depend solely on the clarity of its outward signals, but on its ability to sustain those signals appropriately over extended periods of time.

Much of human-robot interaction research has historically taken place in controlled laboratory or clinical settings. These studies are typically brief—often a single session lasting less than an hour—and involve a single participant interacting with a robot in a supervised, constrained, and isolated manner. Such experiments are well suited for testing feasibility: Can the robot work, at least under these ideal experimental conditions? Yet these conditions differ substantially from the environments in which therapeutic technologies must ultimately operate. Outside the laboratory, novelty fades, expectations evolve, and the value of a system is increasingly judged by whether it continues to provide meaningful support over time.

For autism, truly effective therapy does not unfold within a short, single session; instead, it develops over weeks to years, through exposure to novel, real-world situations that test the ongoing relevance and adaptability of learned strategies. Therefore, the robots we build for autism therapy must function autonomously in dynamic, real-world contexts and maintain relevant support over time. However, each component of that goal highlights both a research gap and the central challenge of this work.

Moving beyond the laboratory to the inherently messy, dynamic, and unpredictable real world means confronting questions that seem deceptively simple. For example, even with state-of-the-art perception systems, questions like "Is my user looking at me?" or "Is it my user or the television speaking?" remain difficult to answer reliably, requiring advances in multimodal perception, representation learning, and robustness to noise (see Figure 3). Beyond accurate detection, long-term autonomy demands that robots can reason about the social context to act appropriately. Deciding, for example, "is it socially appropriate to engage my user right now?" This entails probabilistic reasoning, planning under uncertainty, and contextual inference, all classic challenges in AI and autonomy. Ultimately, once embedded in real-world spaces, robots encounter people and situations that cannot be exhaustively anticipated or preprogrammed. Still, they must improvise appropriate, therapeutically relevant responses in real time, independently and without researcher involvement. This requires generalization beyond training data, fault–tolerant design, sustained hardware and software reliability, and more.

For these reasons, minimalist interaction design should be understood not as a reduction in technical ambition, but as a redistribution of it. The simplicity of the robot's outward behavior can still place greater demands on the underlying system: perception must be robust, autonomy must be sustained, and responses must remain contextually appropriate across extended deployments. In therapeutic robotics, the challenge is therefore not merely to design interactions that are legible and interpretable, but to engineer systems capable of sustaining those interactions in the complex settings of everyday life.

Robots are better understood not as artificial humans, but as interactional scaffolds.

CONCLUSION

Looking back, the googly eyes that filled our house were not sophisticated technologies. They did not perceive attention, generate language, or reason about social context. Yet they embodied a simple design intuition: Social skills develop through repeated, low-stakes opportunities for engagement embedded within everyday life. By turning ordinary objects into playful social partners, my parents transformed our home into an environment rich with invitations to practice eye contact.

Robots extend this same idea, but with far greater flexibility. They can notice where a person is looking, pause to invite a response, repeat an interaction without fatigue, and adapt their behavior across many moments of daily life. In doing so, they can create structured opportunities for practice that would be difficult for human partners to sustain alone.

As robotics and AI continue to advance toward increasingly general and capable systems, it is tempting to measure success by how closely machines resemble our human qualities. Yet in autism therapy, effectiveness does not necessarily follow from realism. The most useful robots may instead be those that communicate clear boundaries about what they can do, provide simple and interpretable signals, and embed social practice within the rhythms of everyday life.

Seen from this perspective, the future of robots in autism interventions may not depend on building ever more human-like machines. Rather, it lies in carefully designing systems that create meaningful opportunities for engagement in the real world. In that sense, the lesson of the googly eyes still holds—sometimes the most powerful interventions are those that make social interaction just a little easier to begin.

References

[1] Scassellati, B. et al. Improving social skills in children with ASD using a long-term, in-home social robot. *Science Robotics* 3, 21(2018.)

[2] Ramnauth, R., Shic, F., and B. Scassellati, B. Gaze behavior during a long-term, in-home, social robot intervention for children with ASD. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction*. IEEE (2025), 949–957.

[3] Scassellati, B., Admoni, H., and Matarić, M. Robots for use in autism research. *Annual Review of Biomedical Engineering* 14 (2012), 275–294.

[4] Pennisi, P. et al. Autism and social robotics: A systematic review. *Autism Research* 9, 2 (2016), 165–183.

[5] Chun, A. et al. Breathe easy: Harnessing robots for stress reduction during pediatric oral challenges. *Journal of Allergy and Clinical Immunology* 155, 2 (2025).

[6] J.-Y. Sung, et al. "My Roomba is Rambo": Intimate home appliances. In Krumm, J., Abowd, G.D., Seneviratne, A., Strang, T. (eds.) *UbiComp 2007: Ubiquitous Computing*. Lecture Notes in Computer Science, vol 4717. Springer (2007), 145–162.

[7] Salomons, N. et al. The impact of an in-home colocated robotic coach in helping people make fewer exercise mistakes. In *Proceedings of the 2022 IEEE International Conference on Robot and Human Interactive Communication*. IEEE (2022), 149–154.

[8] Tsoi, N. et al. Challenges deploying robots during a pandemic: An effort to fight social isolation among children. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*. ACM (2021), 234–242.

[9] Ramnauth, R. et al. A social robot for improving interruptions tolerance and employability in adults with ASD. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction*. IEEE (2022), 4–13.

[10] Ramnauth, R., Bršĉić, D., and Scassellati, B. A robot-assisted approach to small talk training for adults with ASD. In *Proceedings of the Robotics: Science and Systems*. RSS Foundation (2025).

[11] Ramnauth, R., Bršĉić, D., and Scassellati, B. From fidgeting to focused: Developing robot-enhanced social-emotional therapy (RESET) for school de-escalation rooms. In *Proceedings of the 2025 IEEE International Conference on Robot and Human Interactive Communication*. IEEE (2025), 966–973.

Biography

Rebecca Ramnauth is a postdoctoral fellow in computer science at Yale University. Her research explores how interactive AI systems shape human behavior in real-world environments, combining empirical human-AI interaction studies with the development of design principles for safe and reliable autonomy. Her doctoral work focused on long-term deployments of social robots supporting social and emotional development in autism interventions.