

From AI Companionship to Social Scaffolding: A Supervised Avatar Model for Skill Generalization in Autistic Children.

Abstract:

Autistic children receiving applied behavior analytic services may demonstrate a skill in a teaching session without using it with new people or settings. This conceptual article reviews studies published since 2021 on AI-supported autism intervention and specifies a supervised artificial intelligence (AI) avatar for individualized rehearsal between registered behavior technician (RBT) sessions. Early conversational studies report transfer to human dialogue in autistic adolescents and adults, while broader research uses heterogeneous robots, games, virtual reality, and clinician-facing platforms. These findings do not establish the effectiveness of a child-facing avatar embedded in RBT care. The proposed model defines eligible skills, clinician approval, child choice, variation and fading, scheduled human practice, and stopping rules. Its evaluation separates avatar performance from unprompted use with unfamiliar people and materials, and measures acceptability, adverse events, fidelity, and staff burden. Participatory feasibility work followed by controlled comparisons should determine whether the supervised avatar adds practical value beyond equivalent human-supported practice.

Keywords: autism; artificial intelligence; behavioral skills training; generalization; digital intervention; caregiver; implementation

Introduction

A child may learn to ask for help from a familiar technician at a table and still struggle to ask a different adult in a noisy classroom. The clinically meaningful question is whether a skill survives changes in people, materials, cues, and settings. An assessment across home, school, and an unfamiliar research setting illustrates why generalization must be measured explicitly (Carruthers et al., 2025). Repeated digital performance alone therefore cannot establish real-world benefit.

This paper proposes a constrained AI avatar as a supervised rehearsal partner within an existing individualized intervention. It may provide additional, varied practice between visits when a child would benefit from repetition in a different context. This access rationale is a hypothesis to test, not an established solution to staffing or scheduling constraints. Conversational AI trials in autistic adolescents and adults have reported gains that transfer to human dialogue on specified skills, while virtual-reality studies provide separate evidence for technology-assisted practice (Gayle et al., 2025; Koegel et al., 2026a, 2026b; Mittal et al., 2024). These results motivate a child-focused test; neither staffing substitution nor effectiveness of the proposed supervised model follows from them.

The proposal extends Beckman's (2026) published account of synthetic attachment displacement. That earlier work centered relational risk when synthetic companions occupy the place of human support. Here the distinct question is whether a time-limited, clinician-governed simulation can help a child practice a named skill and then use it with actual people. The unit of benefit is observed human-context participation, not affection for or engagement with the avatar. This is a proposed model and research agenda; no efficacy or safety results are claimed.

AI and avatar interventions in autism research since 2021

Scope of the evidence

Since 2021, studies have examined conversational agents, AI-mediated games, social robots, AI-assisted delivery of behavioral procedures, and platforms combined with ongoing therapy. These categories cannot be treated as interchangeable tests of a conversational avatar. A 2025 scoping review of generative AI for autism included 10 studies spanning diagnosis, assessment, intervention, and caregiver information; it did not identify 10 trials of avatar treatment (Sohn et al., 2025). A 2026 systematic review included 14 studies of AI-supported therapeutic interventions across robotics, virtual reality, and wearables (Kuca et al., 2026). A separate meta-analysis included 14 AI-driven game publications, reporting heterogeneity and limited available evidence (Lu et al., 2026). The review totals may overlap and are not a census of distinct child-avatar interventions. They show an emerging field with limited direct applicability to the present model.

Conversational practice and transfer

Koegel et al. (2026b) randomly assigned verbally conversational autistic participants aged 11–35 years to a large-language-model chatbot or waitlist. After four weeks of brief repeated practice, the intervention group improved targeted empathetic statements in human conversation probes. This is direct evidence that one constrained AI conversation task can transfer beyond the interface in the studied population. The comparison with a waitlist does not isolate an AI-specific benefit from additional structured practice; the trained goal was narrow, and the sample did not establish effectiveness for younger clients or children using other communication modes. The study also does not establish a system for ongoing clinician review, RBT-supported home practice, or diverse functional targets.

In a separate 2026 study, Koegel, Ponder, Bruzzese, and Lam reported improved question asking following AI-guided rehearsal, with gains observed in human conversation and maintained over the study's follow-up. This strengthens the rationale for probing human transfer rather than assuming it cannot occur. It remains an initial test of a specified conversational response in adolescents and adults, rather than evidence that an open-ended avatar will generalize across children, skills, or settings (Koegel et al., 2026a).

A conversational virtual-human study with autistic adults explored willingness to interact and perceived social connectedness, but its usability and relationship questions differed from testing an RBT-linked treatment goal or child-level functional transfer (Xyngkou et al., 2024). Taken together, these studies justify a constrained rehearsal hypothesis and accessible, opt-in interaction; they cannot support broad claims that avatar companionship is therapeutic.

AI-mediated child interventions and clinical delivery

AI-driven games and robot-assisted interventions offer varied prompts and structured interaction, but their embodied or gamified formats, target behaviors, and outcome measures differ from dialogue with a screen-based avatar (Kuca et al., 2026; Lu et al., 2026). In an observational study of 43 children aged 2–18, an AI-based platform was used alongside continuous therapy over 12 months. Without a randomized or otherwise comparable group receiving the same therapy without the platform, changes cannot be attributed specifically to AI; neither can findings be assumed to generalize to the present supervised-avatar workflow (Atturu et al., 2025).

A small implementation study used AI guidance to support behavioral-intervention fidelity among two technicians in training and three caregivers; most, but not all, users improved (Yagafarova et al., 2025). Its outcome was implementation by adults, not a child's rehearsal or

independent real-world use. The study nevertheless supports measuring RBT and caregiver fidelity rather than assuming that remote instructions are applied consistently.

What remains to be tested

The relevant gap is a package and population question: whether a clinician-approved conversational avatar used between sessions by autistic children adds measurable transfer across human partners and natural settings when embedded in an individualized, RBT-supported program. Published adolescent and adult chatbot studies provide positive precedents for specific conversational skills, including human transfer (Koegel et al., 2026a, 2026b). Evidence is thinner on younger or minimally speaking children, nonverbal and augmentative communication access, long-term maintenance, comparison with matched human-led rehearsal, child-defined acceptability, privacy and adverse events, and the staff time needed for safe supervision. This conceptual model supplies a protocol and falsifiable outcomes for those unanswered questions; it is not presented as an established intervention.

Mechanism and Boundaries

The hypothesized pathway has four links: clinician-selected target, low-stakes rehearsal with varied cues, prompt fading and feedback, and opportunities with caregivers, technicians, peers, or community partners. Failure at any link can interrupt transfer. Practice may increase fluency in a narrow scenario without changing behavior elsewhere. Carruthers et al. (2025) motivate direct observation across settings; Shingleton-Smith et al. (2024) show that caregiver training can be delivered through telehealth with attention to general-case exemplars; and Gayle et al. (2025) offer a bounded precedent for virtual practice of safety and social communication. These are supporting analogies, not demonstrations of the present avatar's effectiveness.

The avatar is an instrument for instruction, not a friend, therapist, authority figure, or emergency contact. Its language should identify itself as a simulation. It cannot promise secrecy, cultivate exclusive attachment, issue clinical advice, improvise new behavioral goals, or reward continued conversation as an end in and of itself. If a child seeks comfort, shows distress, or introduces a situation outside the approved scenario, the system stops the exercise and directs the child to a designated human. The displacement risk described by Beckman (2026) makes relational boundaries an outcome to monitor, not merely a disclaimer.

Autistic children and families differ in communication style, sensory preferences, interests, and desired outcomes. A child's refusal or nonresponse is meaningful feedback, not a behavior to extinguish for the sake of completing digital sessions. Autistic adults have identified differences between goals that support quality of life and goals aimed at normalization, making social validity of the selected target a substantive design question (Baiden et al., 2025; Schuck et al., 2024). Stakeholder-led priorities in autism intervention (Walton et al., 2024) and participatory digital-design research (Maun et al., 2024) further justify involving autistic participants and caregivers in scenario choice, accessible controls, and acceptability criteria before trial deployment.

Intervention specification

Eligibility and target selection.

A supervising behavior analyst, the child to the extent accessible, and the family select one functional goal already present in the treatment plan. Candidate goals include requesting clarification, asking for a break, negotiating a turn, recognizing a cue to seek help, and practicing a planned response to a routine change. Selection asks whose need the goal serves and avoids targets designed primarily to suppress harmless autistic expression or require a preferred

communication mode. Suitable targets have an observable response, a plausible safe simulation, and a planned real-person opportunity. The team records the child's preferred communication mode, baseline performance, prompts required, sensory access needs, and the contexts in which the skill matters. Rehearsal is optional and stops on child request (Baiden et al., 2025; Schuck et al., 2024).

Targets involving high-risk elopement, aggression, self-injury, abuse disclosure, trauma processing, or emergency judgment require direct professional procedures and are excluded from autonomous avatar rehearsal. Requests for safety or emotional support trigger a preapproved contact pathway. An avatar should not imitate a real classmate, clinician, or caregiver without informed permission; synthetic peers should be visibly fictional.

A repeatable session.

The clinician approves a small scenario library, allowed response range, prompt hierarchy, feedback wording, maximum session length, and escalation rules. The RBT introduces and models a scenario during ordinary care. At home, the child can begin a brief rehearsal with caregiver awareness; the avatar states the goal, plays a plausible partner, waits for the child's response, provides the least intrusive approved prompt, and ends with a concrete invitation to try the skill with a person. Variants change one relevant feature at a time: speaker, wording, visual cues, context, or timing. The system records selected scenarios, prompts, child stop signals, and deviations for clinical review, subject to minimum-data retention.

At the next human session the RBT samples the skill with different materials and a real partner, without copying the avatar's exact wording. The clinician reviews objective probes and adjusts or ends the avatar assignment. A successful simulated exchange never automatically advances a treatment goal. AI may generate candidate phrasings within fixed boundaries, but a

human authorizes scenario families and can suspend the feature. Koegel et al. (2026b) offer a related example of AI-assisted practice of empathetic statements; that study does not validate this proposed supervision architecture.

A supervised path from practice to human use

Each stage names the responsible person and the evidence needed to continue.

STAGE	OWNER AND ACTION	EVIDENCE RECORDED
01 SELECT	Clinician + family Define one meaningful skill and real-world contexts	Baseline; child preference; approved scenarios
02 MODEL	RBT Demonstrate and practice the skill with the child	Prompt level; response definition
03 REHEARSE	Child + avatar Opt into short variants with approved feedback	Completion; prompts; stop requests; errors
04 TRANSFER	RBT + caregiver Arrange opportunities with real people	Unprompted probes with novel people and materials
05 REVIEW	Clinician Inspect transfer, burden, and safety; revise or stop	Decision and reason

HUMAN OVERSIGHT AT EVERY STAGE

Figure 1: Supervised avatar cycle. Progression depends on child choice, documented evidence, and human review; avatar rehearsal is followed by practice with people.

Handback and stopping rules.

A human handback is scheduled after every small block of rehearsal. Expressed distress, repeated failure despite the approved prompt sequence, requests for personal advice, unusual dependence on the avatar, disclosures of harm, or any safety trigger end the exercise immediately. The system gives a brief, nonjudgmental message directing the child to the designated adult and alerts that adult through the approved channel. If a trigger is ambiguous, the

exercise pauses for human review; the system does not diagnose risk or continue a sensitive conversation. Before home use, the team identifies an available adult, a backup contact, hours of use, and the ordinary emergency pathway. If no responsible adult is available, the session cannot start. A caregiver receives a plain-language prompt for a planned real-person practice opportunity, not an obligation to deliver clinical treatment. If no such opportunity exists, further avatar repetition pauses. Detection accuracy, missed triggers, false alarms, response time, and actual handback completion are feasibility outcomes; proposed thresholds must be piloted and preregistered.

Outcomes and Research Design

The targeted outcome is generalization: the proportion of prespecified opportunities in which the child independently uses the selected functional response with an untrained human partner, untrained materials, or a setting distinct from avatar rehearsal. For each child and target, investigators specify before enrollment what constitutes an opportunity, the natural cue that makes the response useful, the response form accepted across speech, gesture, or augmentative communication, a feasible response window, and what counts as assistance.

Independence excludes avatar-specific wording, prompts from the observer, and a rehearsed script elicited by the partner; an ordinary contextual cue, such as an ambiguous instruction, is allowed. An observer who did not deliver rehearsal scores each opportunity from direct observation or consented recording, with a second observer independently coding a prespecified subset for agreement. For a clarification target, a new adult gives an ambiguous instruction in a classroom or community task, and the child asks for clarification within the prespecified window without a teaching prompt. Repeating the practiced script to the avatar is a process measure. Baseline, intervention, and follow-up probes sample several partners and

195 untrained examples; their schedules and success criteria are specified in advance. Withdrawal is
196 not required when removing an acquired functional skill or support would be inappropriate.

Measuring outcomes beyond the avatar

Real-world skill use is primary; the other domains test delivery, experience, and risk.

DOMAIN	PROPOSED MEASURE	INTERPRETATION GUARDRAIL
TRANSFER PRIMARY	Percent of unprompted success with new people/materials; 4–8 week maintenance	Use masked coding where feasible; report opportunity and context.
PRACTICE AND FIDELITY	Avatar trials, prompt level, scenario variation; RBT and clinician adherence	Avatar scores show practice, not real-world transfer.
CHILD EXPERIENCE	Accessible choice/assent; stop requests; brief preference interview	A child may stop without penalty.
BURDEN AND ACCESS	Caregiver/staff minutes; missed visits; travel and setup time	Greater availability alone does not establish better care.
SAFETY AND RELATIONSHIP	Distress, unsafe output, dependency signals, privacy incidents, human handbacks	Record adverse events and the human response.

SUCCESS REQUIRES TRANSFER WITH PEOPLE

197
198 *Figure 2: Candidate outcome domains and interpretation guardrails. The proposed measures*
199 *should be finalized with autistic children, families, clinicians, and the study team.*

200 The first phase should co-design interface and scenarios with autistic children, caregivers,
201 RBTs, and supervisors, followed by a small feasibility study assessing access, acceptability, and
202 failure modes (Maun et al., 2024; Walton et al., 2024). A subsequent concurrent multiple-
203 baseline-across-participants study would stagger the introduction of supervised avatar practice to
204 ongoing care across at least three independently measured participant tiers. Investigators would
205 prespecify allowable baseline lengths, assign intervention start points randomly from those

allowable times, collect repeated natural-context probes in every tier during each phase, and document concurrent changes in ordinary care(Levin & Ferron, 2021).

Replication of change after staggered introduction, observer agreement, fidelity, and prespecified visual analysis would be the principal evidence; a randomization-based analysis could supplement it (Levin & Ferron, 2021). This design tests the effect of adding the whole supervised-avatar package to ongoing care. If feasibility and initial effects justify a later trial, equalize practice time and human supervision and compare avatar rehearsal with matched human-led or nonavatar practice. Only that contrast can begin to isolate the avatar’s added value from extra repetitions.

Potential moderators include language mode, age, sensory access, baseline target skill, caregiver time, prior device use, and setting. Report child-level variation and adverse events, including when the intervention fails. Technology studies can be encouraging yet heterogeneous (Mittal et al., 2024), so pooling practice accuracy with naturalistic transfer would obscure the central question. Qualitative interviews should ask whether the child experiences the tool as useful, burdensome, confusing, or socially displacing.

Implementation, Safety, and Equity

A deployment protocol should identify the supervisor who approves prompts, reviews logs and safety events, corrects unexpected outputs, trains staff, and decides whether use can resume. It should also identify the adult and backup adult available during permitted hours, the notification channel, and how handback completion is documented. Supervisors retain authority to revise goals and determine whether a target remains appropriate. RBTs receive short scripts for modeling and probe delivery, with periodic observation to establish procedural fidelity. Recent work on digital support for procedural fidelity suggests implementation aids are possible

but must be evaluated in the actual workflow (Yagafarova et al., 2025). Possible access constraints motivate testing a low-burden supplement; unobserved AI work must not be counted as equivalent clinical service.

The child-facing system should disclose that responses are generated; restrict outputs to the approved skill domain; avoid diagnoses, treatment changes, coercive reinforcement, and unsolicited emotional intimacy; allow easy exit; and route unexpected or dangerous input to an identified adult. The literature on mental-health conversational agents identifies explainability and unsafe responses as material design problems (Sarkar et al., 2023). For minors, consent from guardians, developmentally appropriate assent, age-appropriate privacy notices, minimized data collection, secure access, retention limits, and incident review should be set before study enrollment. Access includes nonverbal response modes and low-bandwidth alternatives where possible.

Safeguard evaluation should include simulated adversarial inputs and real-world incident monitoring before and during any trial. Investigators should distinguish prototype reliability from clinical effectiveness. A benefit claim should require observed transfer plus acceptable burden and safety; a product that increases screen engagement while reducing human opportunities fails the model's own objective.

Limitations and Alternative Explanations

This article proposes a mechanism and reports no participant data. The reviewed studies cannot be pooled as tests of one treatment: conversational chatbots, games, social robots, virtual reality, and implementation aids differ in interface, dosage, participant age, language ability, comparators, and outcome definition (Kuca et al., 2026; Sohn et al., 2025). Positive human-conversation probes in two chatbot studies warrant replication in children, but do not

demonstrate transfer to multiple natural settings, sustained participation, or outcomes that children and families themselves prioritize (Koegel et al., 2026a, 2026b). Reviews identify small, heterogeneous studies and limited maintenance data (Lu et al., 2026). A caregiver-supported deployment could also select families with more time, connectivity, or comfort with technology. Improvements might reflect extra repetitions, novelty, caregiver attention, concurrent RBT instruction, or maturation rather than the avatar.

The proposed comparison with equal-duration nonavatar practice addresses the additional-practice explanation; prespecified natural-setting probes, varied partners, and follow-up address narrow task learning. Benefits could differ by age, communication mode, sensory access, and the social value of a specific goal. A young child's refusal, confusion, or dependence should be recorded as an outcome, not reclassified as noncompliance. AI output errors and child-data exposure remain possible even under clinician-approved boundaries (Sarkar et al., 2023). Trials must report failures, missingness, supervision burden, and adverse events with the same care as successful rehearsals.

Conclusion

A supervised avatar may offer additional rehearsal of carefully chosen skills between human-led sessions. The decisive test is whether a child uses those skills voluntarily with actual people in meaningful contexts, with acceptable burden and no displacement of human support. A bounded intervention, explicit handback, and independently observed transfer outcomes make that hypothesis testable.

UNDER PEER REVIEW IN IJAR

273 **Declarations**

274 Article type: Conceptual model and research agenda.

275 Ethics approval and consent to participate: Not applicable; no human participants were enrolled.

276 Data availability: No new dataset was generated or analyzed for this article.

277 Funding: None.

278 Competing interests: None.

279 AI assistance: The author used generative AI to assist with reviewing wording and structure,
280 identifying references for verification, and creating the figures. The author reviewed and revised
281 the resulting text, checked the cited sources, and takes responsibility for the manuscript's content
282 and accuracy.

283 Author affiliation and correspondence: The author is an independent scholar.

284

285

References

- Atturu, H., Naraganti, S., & Rao, B. R. (2025). Effectiveness of artificial intelligence-based platform in administering therapies for children with autism spectrum disorder: 12-month observational study. *JMIR Neurotechnology*, 4, e70589. <https://doi.org/10.2196/70589>
- Baiden, K. M. P., Williams, Z. J., Schuck, R. K., Dwyer, P., & Wang, M. (2025). The social validity of behavioral interventions: Seeking input from autistic adults. *Journal of Autism and Developmental Disorders*, 55, 1172–1186. <https://doi.org/10.1007/s10803-024-06297-3>
- Beckman, T. (2026). Synthetic attachment displacement: A developmental framework for adolescent AI companionship. *International Journal of Social Sciences and Humanities Invention*, 13(7), 9031–9043. <https://doi.org/10.18535/ijsshi/v13i07.02>
- Carruthers, S., Charman, T., Leadbitter, K., Ellis, C., Taylor, L., Moore, H., Taylor, C., James, K., Balabanovska, M., Langhorne, S., Aldred, C., Slonims, V., Grahame, V., Howlin, P., McConachie, H., Parr, J., Emsley, R., Le Couteur, A., Green, J., ... Pickles, A. (2025). Generalisation of social communication skills by autistic children during play-based assessments across home, school and an unfamiliar research setting. *Journal of Autism and Developmental Disorders*, 55, 2203–2216. <https://doi.org/10.1007/s10803-024-06370-x>
- Gayle, R. I., Valentino, A. L., & Fuhrman, A. M. (2025). Virtual reality training of safety and social communication skills in children with autism: An examination of acceptability, usability, and generalization. *Behavior Analysis in Practice*, 18, 179–195. <https://doi.org/10.1007/s40617-024-00968-4>
- Koegel, L. K., Ponder, E., Bruzzese, T., & Lam, M. S. (2026a). Using artificial intelligence to improve question asking in autistic adolescents and adults. *Behavior Modification. Advance online publication*. <https://doi.org/10.1177/01454455261477050>
- Koegel, L. K., Ponder, E., Bruzzese, T., Wang, M., Semnani, S. J., Chi, N., Koegel, B. L., Lin, T. Y., Swarnakar, A., & Lam, M. S. (2026b). Using artificial intelligence to improve empathetic statements in autistic adolescents and adults: A randomized clinical trial. *Journal of Autism and Developmental Disorders*, 56, 2513–2529. <https://doi.org/10.1007/s10803-025-06734-x>
- Kuca, J., Stencel, M., Pilarski, B., Florek, S., & Pudlo, R. (2026). Artificial intelligence-supported therapeutic interventions for autism spectrum disorder: A systematic review. *Frontiers in Psychiatry*, 17, 1833853. <https://doi.org/10.3389/fpsyt.2026.1833853>

- 321 Levin, J. R., & Ferron, J. M. (2021). Different randomized multiple-baseline models for different
322 situations: A practical guide for single-case intervention researchers. *Journal of School*
323 *Psychology*, 86, 169–177. <https://doi.org/10.1016/j.jsp.2021.03.003>
- 324 Lu, H., Zhao, J., Zhang, Y., & Cheng, W. (2026). Can AI-driven games enhance social skills for
325 autistic children? A three-level meta-analysis. *Frontiers in Psychology*, 17, 1847426.
326 <https://doi.org/10.3389/fpsyg.2026.1847426>
- 327 Maun, R., Fabri, M., & Trevorrow, P. (2024). Participatory methods to engage autistic people in
328 the design of digital technology: A systematic literature review. *Journal of Autism and*
329 *Developmental Disorders*, 54, 2960–2971. <https://doi.org/10.1007/s10803-023-06015-5>
- 330 Mittal, P., Bhadania, M., Tondak, N., Ajmera, P., Yadav, S., Kukreti, A., Kalra, S., & Ajmera, P.
331 (2024). Effect of immersive virtual reality-based training on cognitive, social, and
332 emotional skills in children and adolescents with autism spectrum disorder: A meta-
333 analysis of randomized controlled trials. *Research in Developmental Disabilities*, 151,
334 104771. <https://doi.org/10.1016/j.ridd.2024.104771>
- 335 Sarkar, S., Gaur, M., Chen, L. K., Garg, M., & Srivastava, B. (2023). A review of the
336 explainability and safety of conversational agents for mental health to identify avenues
337 for improvement. *Frontiers in Artificial Intelligence*, 6, 1229805.
338 <https://doi.org/10.3389/frai.2023.1229805>
- 339 Schuck, R. K., Geng, A., Doss, Y., Lin, F., Crousore, H., Baiden, K. M. P., Dwyer, P., Williams,
340 Z. J., & Wang, M. (2024). A qualitative investigation into autistic adults' perspectives on
341 intervention goals for autistic children. *Neurodiversity*, 2.
342 <https://doi.org/10.1177/27546330241266718>
- 343 Shingleton-Smith, C., Koudys, J., Azzano, A., & Feldman, M. (2024). Telehealth general case
344 parent training for children at risk for autism. *Journal of Applied Behavior Analysis*,
345 57(4), 879–892. <https://doi.org/10.1002/jaba.2913>
- 346 Sohn, J.-S., Lee, E., Kim, J.-J., Oh, H.-K., & Kim, E. (2025). Implementation of generative AI
347 for the assessment and treatment of autism spectrum disorders: A scoping review.
348 *Frontiers in Psychiatry*, 16, 1628216. <https://doi.org/10.3389/fpsyt.2025.1628216>
- 349 Walton, K. M., Borowy, A. R., Gordon, R. A., & Wainer, A. L. (2024). Enhancing stakeholder
350 roles in autism early interventions in the United States: A stakeholder-driven research
351 agenda. *Autism*, 28(5), 1120–1134. <https://doi.org/10.1177/13623613231195743>
- 352 Xygykou, A., Siriaraya, P., She, W.-J., Covaci, A., & Ang, C. S. (2024). “Can I be more social
353 with a chatbot?”: Social connectedness through interactions of autistic adults with a
354 conversational virtual human. *International Journal of Human–Computer Interaction*,
355 40(24), 8937–8954. <https://doi.org/10.1080/10447318.2023.2292880>

356 Yagafarova, A., Dowling, E., LaMonica, T., Hantula, D. A., Nosek, J. T., & Jimenez-Gomez, C.
357 (2025). Supporting procedural fidelity of behavioral interventions for children with
358 autism via an artificial intelligence platform. *Behavioral Interventions*, 40(1), e2075.
359 <https://doi.org/10.1002/bin.2075>

UNDER PEER REVIEW IN IJAR