



RESEARCH ARTICLE

SYNTHETIC THERAPEUTIC SCAFFOLDING FOR SKILL GENERALIZATION IN AUTISTIC CHILDREN

Timothy Beckman

1. DHM, RBT Independent Scholar.

Manuscript Info

Manuscript History

Received: 16 July 2026

Final Accepted: 18 August 2026

Published: September 2026

Key words:-

autism intervention; skill generalization;
conversational agents; AI avatar;
human-AI interaction; behavioral skills
training

Abstract

Autistic children may learn a functional response during instruction without using it when people, settings, or demands change. Synthetic Therapeutic Scaffolding (STS) is proposed as a clinician-governed approach to this problem. An AI avatar would provide brief, restricted rehearsal of an individualized skill, followed by practice with people and gradual removal of synthetic support. A targeted narrative synthesis identifies relevant evidence from conversational agents, virtual reality, and implementation platforms while examining the limits of its application to children. The proposed model connects varied practice and prompt fading with independently observed human transfer. Its requirements include accessible communication, ongoing child assent, controlled outputs, adult intervention, and defined privacy and safety procedures. A concurrent multiple-baseline study would evaluate the supervised package, followed by a comparison with human-led rehearsal matched for time, opportunities, and adult attention. Generalization is the primary outcome; maintenance, acceptability, fidelity, burden, and adverse events require separate assessment. Existing evidence supports investigation of this approach but does not establish its effectiveness or safety. Participatory design, technical testing, and feasibility work must precede efficacy evaluation.

"© 2026 by the Author(s). Published by IJAR under CC BY 4.0. Unrestricted use allowed with credit to the author."

Introduction:-

A child may ask a familiar technician for help at a table but have difficulty making the same request to another adult in a noisy classroom. The response has been learned under one set of conditions; its practical value depends on whether it remains available when those conditions change. Assessment across home, school, and unfamiliar settings shows why generalization must be measured directly (Carruthers et al., 2025). Performance within a digital exercise cannot establish that a child can use the skill with people.

Synthetic Therapeutic Scaffolding (STS) would provide temporary rehearsal within an individualized program governed by a clinician and supported by a registered behavior technician (RBT). The avatar would practice one approved skill, then direct practice back to human partners as synthetic support recedes. The intent is to increase the child's capacity to participate in everyday situations. This proposal extends the concern about relational displacement described in synthetic attachment displacement (Beckman, 2026), but it does not assume that brief

digital practice causes dependence or that supervision prevents it. No participant data, working prototype, or validated protection against dependence is reported. The proposed contribution is a framework that connects functional target selection, accessible child choice, restricted rehearsal, a defined return to human practice, and independent measurement of transfer. Related technologies already provide conversational practice and clinician control. STS brings these elements together around an explicit end state: voluntary skill use with people after the synthetic support has been reduced or removed. Whether this integration improves access, reduces costs, or saves staff time remains to be tested.

Evidence identification and selection

The evidence is presented as a targeted narrative synthesis. The submitted draft did not retain a reproducible database search log, duplicate-screening record, or study flow count, so an exhaustive review cannot be claimed. Source verification on October 4, 2026 used web-indexed PubMed and PubMed Central records and journal publisher pages. Exact titles and DOI strings from the draft were searched, supplemented by combinations of autism, artificial intelligence, chatbot, conversational agent, virtual human, virtual reality, generalization, behavioral skills training, and procedural fidelity. A further targeted audit checked the principal conversational studies, observational platform study, and behavioral skills training source on October 4, 2026.

The synthesis considers English-language peer-reviewed work published from January 2021 through the initial verification date. Empirical autism studies and evidence syntheses were selected when they addressed rehearsal, outcomes with people, implementation, or participatory design. Adult studies provide indirect evidence and are identified as such because their findings cannot establish suitability for children. Methodological work and the author's conceptual article inform the framework rather than its efficacy.

Commercial claims and diagnostic-only studies were not used to support treatment effectiveness. Unverifiable references and findings available only through secondary commentary were excluded from efficacy reasoning. Publisher or indexed records were used to check publication details, populations, designs, and claims. Selection and verification were conducted without independent duplication. The resulting account may omit relevant studies and does not provide a pooled effect estimate or formal risk-of-bias assessment; unavailable details are identified rather than inferred.

Recent evidence and its limits:-

The review literature covers substantially different uses of AI. Sohn et al. (2025) examine generative AI in assessment, intervention, and caregiver support; Kuca et al. (2026) review broader AI-supported interventions; and Lu et al. (2026) synthesize AI-driven games. These reviews help identify possible applications, but their findings cannot be treated as independent tests of STS. Virtual-reality trials offer a related rationale for technology-assisted learning while using different interfaces, teaching conditions, and outcomes (Mittal et al., 2024). The technology category alone does not identify the mechanism responsible for improvement.

Targeted conversational rehearsal with Noora provides the most direct precedent. Koegel, Ponder, Bruzzese, Wang, et al. (2026) evaluated empathetic responding in 30 autistic participants aged 11–35 who could engage in conversation. The four-week randomized waitlist comparison included human conversation probes. Koegel, Ponder, Bruzzese, and Lam (2026) examined question asking in four adolescents and adults, using three concurrent multiple-baseline tiers and one nonconcurrent replication, with follow-up assessment of human conversation. These findings support further investigation of rehearsal and transfer. They do not establish effectiveness for younger children or AAC users, and neither study separates AI practice from an equal-duration human-led alternative.

The remaining studies inform different parts of the proposal. Xygkou et al. (2024) examine adults' experiences with a virtual human and provide evidence about usability and social experience. Atturu et al. (2025) followed 43 children using an AI platform alongside therapy for 12 months; adherence-defined groups leave selection and concurrent care as plausible explanations for differences. Yagafarova et al. (2025) assess adult implementers' procedural fidelity, which cannot establish child-avatar transfer. Gayle et al. (2025) examine virtual-reality training and natural-environment probes in three children. Tables 1 and 2 make these distinctions explicit so that evidence for one component is not presented as validation of the complete model.

Table 1: Selected empirical studies informing the model

Study	Participants	Technology and target	Exposure	Comparator
Koegel et al. (2026), question asking	4 adolescents/adults; conversational	Noora LLM chatbot; question asking	20 trials/day, 5 days/week; criterion-based exposure	Concurrent baseline in 3; fourth nonconcurrent
Koegel et al. (2026), empathetic statements	30 participants, ages 11–35; sentence-level conversation	Noora LLM chatbot; empathetic responses	10 trials/day, 5 days/week, 4 weeks	Randomized waitlist
Xygkou et al. (2024)	6 autistic and 6 non-autistic adults	Kuki conversational virtual human; social experience	1–4 weeks	Exploratory group comparison
Atturu et al. (2025)	43 autistic children, ages 2–18	CognitiveBotics plus therapy; developmental skills	12 months	Adherence-defined groups
Yagafarova et al. (2025)	2 trainee technicians; 3 caregivers	GAINS AI platform; intervention fidelity	Phase-based; exact dose not extracted	Within-person implementation phases
Gayle et al. (2025)	3 autistic children, ages 5–8	VR treatment package; safety/social communication	Individualized sessions; exact dose not extracted	Within-participant baseline/treatment

Table 2: Measures and limits of interpretation


Study	Outcome measures	Generalization evidence	Major limitation
Koegel et al. (2026), question asking	AI accuracy; human question asking; confidence and satisfaction	Human probes; 5-week follow-up	Small sample; replacement; no active comparator
Koegel et al. (2026), empathetic statements	Human empathetic response rate; AI accuracy; self-report	Pre/post human conversation	Waitlist; narrow response; limited communication range
Xygkou et al. (2024)	Interviews, chatlogs, questionnaires	No functional transfer test	Small adult qualitative sample
Atturu et al. (2025)	CARS, social maturity, developmental and language scales	No isolated avatar-to-human transfer endpoint	Adherence selection; combined therapies
Yagafarova et al. (2025)	Observed adult procedural fidelity	Implementer performance, not child transfer	Five implementers; one caregiver did not improve
Gayle et al. (2025)	Skill accuracy, maintenance, happiness, social validity	Natural-environment generalization assessments	Small VR package study; no avatar-specific contrast

Note. LLM = large language model; VR = virtual reality; CARS = Childhood Autism Rating Scale. Exact dose not extracted indicates that a numeric exposure was not established in this targeted extraction; it does not imply that the original article omitted it. These are selected exemplars, not an exhaustive study census.

Theoretical mechanism and testable propositions:-

Behavioral skills training combines instruction, modeling, rehearsal, and feedback. Evidence from training professionals supports this instructional structure, although it does not validate replacing a professional with AI (Slane & Lieberman-Betz, 2021). Under STS, the clinician would define the functional response, the RBT would model it, and the avatar would provide approved practice and feedback. General-case training would vary relevant examples, while prompt fading would shift control toward the natural cue. Telehealth general-case parent training offers a related rationale for choosing varied examples (Shingleton-Smith et al., 2024). The intended outcome is independent use after rehearsal ends.

The mechanism requires more than repeated avatar practice. Clinician-approved goals, accessible controls, an available adult, and scheduled human opportunities must support modeling, variation, prompt fading, and a return to practice with people. Reduced prompt dependence and practice fluency are intermediate outcomes. Independent responses in untrained human situations and later maintenance are the outcomes that would establish practical value. A child may instead memorize a phrase, dislike the interface, or have too few human opportunities to use the response. Figure 1 shows these points of failure as part of the proposed model.



Note. All elements describe proposed requirements. Unsafe output, distress, loss of access, or unsuccessful transfer requires human review. Rehearsal success does not authorize progression.

Proposition 1. If the supervised STS package supports generalization, independent responding with people should increase after its introduction at different times across concurrent participants. Replicated changes would support the package; persistent rehearsal gains without human transfer would weaken its central claim.

Proposition 2. If the avatar adds value, transfer should exceed that achieved through human-led rehearsal when time, target opportunities, and adult attention are matched. Equivalent outcomes would weaken an avatar-specific efficacy claim and shift the decision toward burden, access, and cost.

Proposition 3. Varying relevant cues and fading prompts, followed by timely human practice, should produce greater transfer than rehearsal that does not include those features. This mechanism requires a comparison that manipulates the proposed features; improvement in the complete package alone cannot identify their individual contributions.

Proposition 4. Useful implementation should combine child-reported usefulness, completed returns to human practice, and manageable burden without concerning distress or displacement. Favorable skill scores alongside harm or unacceptable burden would not support expansion. These propositions are hypotheses rather than findings.

Population and target selection:-

An initial feasibility study would involve autistic children aged 8–12 who already receive clinician-supervised behavioral services. A functional target would be chosen with each child and family. This age range is a practical starting point, not an established treatment indication. Participation requires an accessible means of understanding the exercise, expressing assent or refusal, and responding through speech, text, gesture, symbol selection, or augmentative and alternative communication (AAC). IQ, eye contact, and fluent speech would not determine eligibility. A communication specialist must test AAC access before enrollment. If the interface cannot support a child's communication, it requires redesign or deferral; the access problem should not be attributed to the child.

The clinician would document accepted responses, processing time, sensory preferences, support needs, and baseline performance. Symbol input, device compatibility, adjustable wait times, visual schedules, and trained interpretation of gestures may be needed for AAC access. Speech recognition must not override the child's communication device. Functional variations in a response would be accepted without requiring the avatar's exact wording, and prompted responses would be recorded separately. This distinction protects the measurement of independence.

Targets could include requesting clarification or a break, asking for help, and negotiating a turn. Their purpose is to support participation, not suppress harmless autistic expression. The supervising clinician determines clinical appropriateness and safety, the guardian provides permission, and the child retains the right to decline or stop. Any involved party, including the RBT, may suspend use. Unresolved disagreement prevents enrollment or pauses the target until discussion and documented review resolve it. Adult agreement cannot override child refusal. Autistic adults' perspectives on intervention goals inform these boundaries (Baiden et al., 2025; Schuck et al., 2024). High-risk behavior management, trauma processing, abuse assessment, and emergency judgment would remain outside rehearsal.

Proposed technical architecture and workflow:-

The first prototype would use a version-controlled scenario engine and a closed library of clinician-approved prompts and feedback. A fixed state machine would control entry, scenario selection, response collection, feedback, and exit. An optional large language model (LLM) could classify responses against an approved rubric or propose restricted wording, but it could not create goals, change treatment, or conduct open-topic conversation. Model identity, version, templates, settings, scenario hashes, and validator versions would be recorded before testing. A closed-library prototype can test the supervised rehearsal process; a separate comparison would be needed to establish whether generative AI contributes anything beyond that process. No vendor or implemented system is claimed here.

Child input would be processed as response data and could not change system instructions. Model output would require schema, domain, and safety validation before display. Invalid, uncertain, or out-of-domain output would be discarded and replaced by a fixed pause message and transfer to the available adult. The system would have no unrestricted browsing, independent contact functions, or long-term relationship memory. An adult dashboard would record deviations and allow immediate suspension. Human review of the approved library, review of flagged events before resumption, and renewed bench testing after updates would be required. Output restrictions reduce exposure; they cannot guarantee safety.

The distinctions between technologies matter because they change what is being tested. A conversational chatbot may permit open topics, whereas STS would restrict practice to one approved objective. A virtual human describes the interface; it does not necessarily define clinical oversight. Virtual reality changes immersion and spatial context, a robot adds physical embodiment, and clinician-facing platforms support adult implementation. STS could use some of these components, but its proposed benefit depends on the supervised process and transfer to people. A comparison that changes only the visible interface would be needed to evaluate the contribution of embodiment.

A session would begin with accessible assent and confirmation that a responsible adult is present and able to intervene. The avatar would identify itself as a practice program and name the skill. An initial ceiling of 5–10 minutes, up to three sessions weekly, is proposed for feasibility testing. The clinician would specify response time, prompts, feedback, and a limit of two unsuccessful trials before returning practice to the adult. Relevant cues would vary, and prompts would decrease according to predefined criteria. The child could leave at any point. Each session would be followed by a planned human practice opportunity, ideally within 24 hours; further avatar practice would pause if that opportunity could not occur. Independent human probes would guide clinical review.

Operational safety and human handback:-

The avatar would use a visibly fictional identity and could not claim feelings, friendship, secrecy, exclusivity, or therapeutic authority. Engagement streaks, rewards for extended use, and stored relationship narratives would be excluded. Simple examples and accessible reminders would discourage identifying disclosures, while an adult would remain available to respond. These controls are intended to limit relational pressure and disclosure exposure. Their presence does not establish that the child understands the boundary or that automated detection will identify every concern.

A stop signal, reported discomfort, or individually agreed observable sign such as crying or attempting to leave would end rehearsal immediately. Refusal would not be treated as an error. Possible dependence would be considered when attempts to bypass limits, marked distress at termination, or reduced participation in previously valued human activities occurred relative to the child's baseline and context. Two observations within seven days would trigger clinical review; one serious event would require immediate suspension. These proposed thresholds identify a need for review and do not diagnose dependence.

Unsafe output would include personal or clinical advice, secrecy requests, coercion, discriminatory content, fabricated authority, or an out-of-domain response. Such output, a harm disclosure, an ambiguous safety cue, repeated unsuccessful rehearsal, or a system failure would pause the session. The interface would direct the child to the adult who is already present. The adult would acknowledge the alert immediately and confirm contact within two minutes. An unacknowledged alert would keep the interface locked and display the approved backup contact route. These timings are proposed service targets for testing, not a reason to delay immediate intervention. The adult would follow the existing safety or emergency plan; the program would not assess or manage a crisis.

Bench testing would include labeled examples of refusal, indirect disclosure, AAC expressions, prompt injection, harmful input, out-of-domain requests, network loss, and notification failure. Sensitivity, missed safety cues, false alarms, blocked unsafe outputs, validator bypasses, and alert and adult-contact latency would be reported separately for each language and communication mode. An unblocked severe response or missed severe test cue would prevent child deployment until correction and retesting. A test set with no failures cannot establish zero risk. Trial records would identify event time, context, severity, system version, action, responsible adult, resolution, and recurrence. Serious events would stop use pending independent safety review; minor false alarms would remain part of the burden assessment.

Ethics and data governance:-

An empirical study would require prospective ethics review, guardian consent, accessible ongoing child assent, and clear withdrawal procedures. Children and guardians would receive an explanation of the program, its records, access to those records, and the response to disclosures. Declining or withdrawing could not affect usual care. The investigator and supervising clinician would remain accountable, with named technical and data leads. Reliability must be tested because a disclaimer cannot address the documented safety problems of conversational agents (Sarkar et al., 2023).

Store a random study identifier, scenario/version identifiers, response category, prompt level, timestamps, and event flags. Keep the identity key separately. Raw child audio/video is off by default; any necessary recording needs separate consent and a defined purpose. Avoid retaining free-text conversation; redact incidental identifiers before any approved external processing. Sensitive disclosures are transferred minimally to the designated professional record and are not kept in routine rehearsal logs or used for model training.

A proposed study schedule deletes transient response content within 24 hours and separately consented probe recordings within 30 days after adjudicated coding. Retain deidentified coded outcomes and event summaries for five years after study closure, unless an approved institutional requirement specifies another period. The consent

form must state the actual schedule before enrollment. Withdrawal stops collection; delete identifiable research records when permitted and explain exceptions for already anonymized results or required incident records. Verify deletion across backups under the approved backup-expiry schedule, and document completion.

Require encryption in transit and at rest, role-based access, multifactor authentication for staff, access logs, separate identity storage, and tested recovery and incident response. RBTs access assigned targets and needed flags; supervisors access clinical events; independent observers access pseudonymized probes; technical staff receive redacted diagnostics. Vendor contracts must prohibit training and secondary use, establish retention and processing location, and permit deletion. No compliance status is inferred from a vendor label. A breach pauses collection and follows the approved institutional notification and response plan. Review hallucinations, dialect bias, AAC misclassification, and uneven safety performance with autistic stakeholders before release.

Evaluation protocol and outcomes:-**Primary endpoint and measurement:-**

The primary endpoint would be the proportion of prespecified opportunities with people in which the child produces an independent functional response. Results would be summarized across the final two intervention weeks relative to baseline and plotted as a repeated series for each child. Five dimensions require separate measurement: people who did not teach the target, different partner roles such as a peer instead of an adult, settings outside rehearsal, untrained materials, and changed demands such as wording, timing, or routine. Reports would distinguish each dimension and their combinations. Success under one change does not demonstrate transfer under the others.

Before baseline, the protocol would define the natural cue, a valid opportunity, accepted response forms, accessible response time, and any teaching prompts that disqualify independence. For example, an untrained adult could give an ambiguous instruction and allow the child to request clarification through speech or AAC. The natural cue would remain available, but instructional prompts would not. An observer independent of rehearsal would administer or code probes, with masking to phase where practicable. The avatar would not score the primary outcome.

Three probe sessions weekly would provide at least five valid opportunities per session where feasible, drawn from a rotating prespecified bank. Final assessment would include at least two untrained partners, two settings, and two material sets. A provisional decision rule would require a gain of at least 20 percentage points from baseline and at least 80% independent responding across three consecutive probes, with demonstration across two untrained partners and two settings. These numbers are illustrative and are not established minimal clinically important differences. Children, families, and clinicians would agree on final thresholds before preregistration and outcome inspection.

Secondary and safety outcomes:-

Secondary outcomes include independent transfer at two and six weeks after rehearsal ends, voluntary skill use reported and corroborated in natural routines, and child-rated usefulness and comfort using an accessible three-choice scale weekly and at exit. Caregiver and staff interviews provide separate perspectives rather than substitutes for assent. Process measures are prompts, rehearsal attempts, and completed human opportunities. Fidelity is completed prescribed steps divided by applicable steps, measured in at least 30% of sessions across staff and phases. Record supervision, preparation, troubleshooting, and caregiver time in minutes, plus device/connectivity costs and missed access. Safety endpoints include events per session and per child, severity, false alarms, missed cues, completed handbacks, and changes in human participation. Figure 2 distinguishes these outcome classes.

Figure 1: Outcome classes and interpretation safeguards

PRIMARY	Independent human-context responding Avatar accuracy is not transfer
SECONDARY	Maintenance at 2 and 6 weeks; child usefulness Report dimensions and child-level variation
PROCESS	Prompt levels, repetitions, human practice, fidelity Exposure cannot establish benefit
IMPLEMENTATION	Staff/caregiver minutes; costs; access failures Measure displaced burden and missingness
SAFETY	Distress, unsafe output, missed cues, handback Benefit requires acceptable risk and participation

Note. The primary endpoint is scored independently of the avatar. Process, implementation, and safety measures remain separate from efficacy outcomes.

A second independent observer codes at least 30% of probe sessions, balanced across participants, phases, partners, and communication modes, plus all consented serious-event recordings. Report opportunity-by-opportunity agreement and chance-corrected agreement where estimable. Proposed training targets are at least 85% agreement and 90% procedural fidelity. Scores below target trigger retraining and review; original disagreements and affected observations remain reported. Resolve discrepancies by prespecified adjudication without hiding pre-adjudication reliability.

Missing observations require explicit accounting. A session without a valid opportunity would be missing, while a nonresponse during a valid completed opportunity could be scored unsuccessful under the agreed definition. A child's withdrawal would stop the probe and be reported separately. Each missed probe and its reason would be recorded, and graphs would show gaps without carrying earlier scores forward. Reports would include available-opportunity denominators, missingness by phase and cause, and sensitivity bounds that treat missing planned opportunities as unsuccessful or successful. Excluding difficult sessions would conceal a limitation rather than resolve it.

Staged design and experimental control:-

Participatory interface development and bench testing would precede a feasibility cohort of approximately six children with varied communication access and support needs. The purpose of this cohort would be to test workflow, not estimate efficacy. Autistic adults and children would help choose targets, review scenarios, test exit controls, define acceptability, and interpret findings through accessible, compensated activities. A decision log would record which recommendations were adopted or declined and why. This gives participation a defined role in design decisions (Maun et al., 2024; Walton et al., 2024).

A subsequent exploratory concurrent multiple-baseline study would recruit six participants in two replications of three independently measured tiers. Staggered introduction within each replication would allow repeated examination of whether change follows the intervention while concurrent untreated tiers remain unchanged. Six participants would provide replication, not population-level statistical power. The protocol would prespecify baseline windows around 6, 9, and 12 probes and randomize eligible starts within those windows. A proposed readiness rule would require five recent observations without a sustained improving trend and a range no greater

than 20 percentage points. The exact allowable windows, trend criterion, and extension maximum must be registered before enrollment. Any stability-dependent alteration would be reported and incorporated into the analysis; it could invalidate the planned randomization test (Levin & Ferron, 2021).

The intervention phase would provide four weeks of rehearsal, using the same probe schedule as baseline, followed by assessment at two and six weeks. Visual analysis would examine level, trend, variability, immediacy, overlap, and consistency across tiers. Evidence of experimental control would require replicated changes following introduction while untreated concurrent tiers remained unchanged. A prespecified randomization analysis could supplement visual assessment only if the allocation procedure and missingness assumptions remained valid. This design would test the supervised package as a whole; it would not isolate the AI component.

Keep target-related usual care and caregiver practice stable where clinically appropriate; log medication, staffing, school, therapy, and target-teaching changes weekly. Necessary care takes priority over experimental control. Analyze or clearly qualify tiers affected by concurrent changes, contamination, or unstable baselines. A six-participant series cannot determine population efficacy or rare-event safety.

Matched comparative evaluation:-

Further comparison would be justified only if feasibility and initial transfer supported it. Avatar and human-led rehearsal would then be randomized while matching targets, opportunities, session duration, cue variation, feedback, prompts, supervision time, and scheduled human practice. Usual care and independent probes would be equivalent in both groups. Actual practice and attention would be recorded because assigned dose may differ from delivered dose. Pilot variability would inform an a priori sample-size analysis; no powered sample is claimed here. An additional comparison with text-only delivery through the same engine would address embodiment. These contrasts would establish whether the avatar contributes enough benefit to justify its resources.

Implementation equity and limitations:-

Implementation would need to account for device access, connectivity, language, communication supports, and caregiver availability. Loaned devices, connectivity assistance, low-bandwidth controls, culturally reviewed translations, and accessible alternatives should be available. Language and AAC versions require separate testing because results from English speech cannot validate them. Support time, socioeconomic barriers, prior device familiarity, and accommodations would be documented. Families unable to provide continuous supervision would need an adequately staffed human alternative. An intervention cannot be said to improve access if its equipment and supervision requirements create a new barrier.

Families with time, connectivity, or enthusiasm for technology may be more likely to enroll, limiting the population represented. Novelty, observation effects, caregiver involvement, concurrent services, maturation, and baseline communication differences could also account for improvement. Independent assessment and matched comparisons address some of these explanations but cannot eliminate all bias. Recognition and feedback may work unevenly across dialects and communication styles. Short follow-up can examine persistence over the measured period; it cannot establish durable relationships, broad participation, or rare-event safety. Negative outcomes, refusals, burden, and unsafe events must remain visible alongside successful transfer.

Conclusion:-

STS proposes a temporary route from restricted avatar rehearsal to voluntary skill use with people. Its value would depend on transfer that remains after synthetic support recedes, with acceptable safety, access, and burden. The next task is to determine whether the supervised package produces replicated independent responding, whether an avatar adds value beyond matched human practice, and which communication supports make participation possible. Participatory design, access testing, technical reliability, and feasibility work must establish whether controlled comparison is warranted. If transfer is absent or the resources and harms outweigh the benefit, the model should be revised or discontinued.

Declarations:-**Article type:**

Conceptual model and research agenda.

Ethics approval and consent: Not applicable to this article; no participants were enrolled.

Data availability:

No new participant dataset was generated or analyzed.

Funding:-

None.

Competing interests:-

None.

Author affiliation:-

Independent scholar.

AI assistance: Generative AI assisted with manuscript wording and organization, targeted source verification, figure preparation, and the response draft. The author retains responsibility for the argument, reference accuracy, proposed protocol, and final submission.

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. (2026). 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. (2026). 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/fpsy.2026.1833853>
- Levin, J. R., & Ferron, J. M. (2021). Different randomized multiple-baseline models for different situations: A practical guide for single-case intervention researchers. *Journal of School Psychology*, 86, 169–177. <https://doi.org/10.1016/j.jsp.2021.03.003>
- Lu, H., Zhao, J., Zhang, Y., & Cheng, W. (2026). Can AI-driven games enhance social skills for autistic children? A three-level meta-analysis. *Frontiers in Psychology*, 17, 1847426. <https://doi.org/10.3389/fpsyg.2026.1847426>
- Maun, R., Fabri, M., & Trevorrow, P. (2024). Participatory methods to engage autistic people in the design of digital technology: A systematic literature review. *Journal of Autism and Developmental Disorders*, 54, 2960–2971. <https://doi.org/10.1007/s10803-023-06015-5>
- Mittal, P., Bhadania, M., Tondak, N., Ajmera, P., Yadav, S., Kukreti, A., Kalra, S., & Ajmera, P. (2024). Effect of immersive virtual reality-based training on cognitive, social, and emotional skills in children and adolescents

- with autism spectrum disorder: A meta-analysis of randomized controlled trials. *Research in Developmental Disabilities*, 151, 104771. <https://doi.org/10.1016/j.ridd.2024.104771>
13. Sarkar, S., Gaur, M., Chen, L. K., Garg, M., & Srivastava, B. (2023). A review of the explainability and safety of conversational agents for mental health to identify avenues for improvement. *Frontiers in Artificial Intelligence*, 6, 1229805. <https://doi.org/10.3389/frai.2023.1229805>
 14. Schuck, R. K., Geng, A., Doss, Y., Lin, F., Crousore, H., Baiden, K. M. P., Dwyer, P., Williams, Z. J., & Wang, M. (2024). A qualitative investigation into autistic adults' perspectives on intervention goals for autistic children. *Neurodiversity*, 2. <https://doi.org/10.1177/27546330241266718>
 15. Shingleton-Smith, C., Koudys, J., Azzano, A., & Feldman, M. (2024). Telehealth general case parent training for children at risk for autism. *Journal of Applied Behavior Analysis*, 57(4), 879–892. <https://doi.org/10.1002/jaba.2913>
 16. Slane, M., & Lieberman-Betz, R. G. (2021). Using behavioral skills training to teach implementation of behavioral interventions to teachers and other professionals: A systematic review. *Behavioral Interventions*, 36(4), 984–1002. <https://doi.org/10.1002/bin.1828>
 17. Sohn, J.-S., Lee, E., Kim, J.-J., Oh, H.-K., & Kim, E. (2025). Implementation of generative AI for the assessment and treatment of autism spectrum disorders: A scoping review. *Frontiers in Psychiatry*, 16, 1628216. <https://doi.org/10.3389/fpsy.2025.1628216>
 18. Walton, K. M., Borowy, A. R., Gordon, R. A., & Wainer, A. L. (2024). Enhancing stakeholder roles in autism early interventions in the United States: A stakeholder-driven research agenda. *Autism*, 28(5), 1120–1134. <https://doi.org/10.1177/13623613231195743>
 19. Xygykou, A., Siriaraya, P., She, W.-J., Covaci, A., & Ang, C. S. (2024). “Can I be more social with a chatbot?”: Social connectedness through interactions of autistic adults with a conversational virtual human. *International Journal of Human–Computer Interaction*, 40(24), 8937–8954. <https://doi.org/10.1080/10447318.2023.2292880>
 20. Yagafarova, A., Dowling, E., LaMonica, T., Hantula, D. A., Nosek, J. T., & Jimenez-Gomez, C. (2025). Supporting procedural fidelity of behavioral interventions for children with autism via an artificial intelligence platform. *Behavioral Interventions*, 40(1), e2075. <https://doi.org/10.1002/bin.2075>