SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

From Watching to Choosing: A Design Science Account of UD (Understanding Duo)- a Branching-Narrative Animated Curriculum for Social Skills Learning in Autistic Children

[1]Nishant Bhuvanesh Trivedi, [2]Dr. Shruti Tiwari

[1]Vadodara, Parul Institute of Design, Parul University. E-mail: nishant.trivedi89029@paruluniversity.ac.in
[2]Vadodara, Parul Institute of Design, Parul University. E-mail: shruti.tiwari25189@paruluniversity.ac.in

Abstract
Animated instruction and social narratives are both widely used to teach social skills to autistic children, and both are, in their conventional forms, things a child watches or is read to. The child's role is receptive. This paper reports the design of UD (Understanding Duo), an animated social skills curriculum in which each episode **is** instead structured as a decision tree: the child is asked what the character should do next, chooses among three options, and sees that choice play out. We follow the design science research (DSR) paradigm and use the six-activity process model of Peffers and colleagues to structure the account. The designed artefact comprises a curriculum specification of 140 episodes across seven skill levels — foundational social skills, emotional understanding, daily living, school, community, friendship and safety — together with full screenplays, choice architectures and adult support material for each episode. From the design work we abstract seven design principles, formulated using the aim–context–mechanism schema proposed by Gregor, Chandra Kruse and Seidel. The most consequential is a non-punitive branching rule: no choice offered to the child leads to ridicule, failure or an unrecoverable outcome, and consequences differ in usefulness rather than in valence. We situate the contribution in the improvement quadrant of the DSR knowledge contribution framework: the problem — teaching social skills in a way that transfers — is well established; the solution, treating the episode as a navigable decision space rather than a linear text, is not. The paper reports design and demonstration only; no empirical evaluation is claimed.

Keywords: autism spectrum condition; social skills; design science research; branching narrative; animation; design principles; educational technology

1. Introduction

Three decades of intervention research have established, with reasonable confidence, that autistic children can be taught discrete social skills. What has proved much harder is getting those skills to travel. A child who reliably greets a therapist in a clinic room may not greet a classmate in a corridor. This is not a marginal complaint about intervention research; it is arguably its central unresolved problem, and reviewers have repeatedly identified it as the point at which the evidence base thins out (Wang et al., 2011; Qi et al., 2018; Nejati et al., 2024).

Two of the most widely adopted instructional formats share a structural feature that may bear on this. Video modelling shows a child a correct performance to imitate. Social narratives describe a situation and the behaviour expected within it. Both have accumulated respectable support — video modelling in particular has shown large effects in single-case syntheses (Wang et al., 2011) — and both position the child as a recipient. The material specifies the answer; the child's task is to absorb and reproduce it. Where a real social situation is open-ended, the instructional artefact is closed.

There is a plausible argument that this mismatch matters. Social competence in ordinary settings is not the retrieval of a single memorised response but the appraisal of a situation in which several responses are available and none is uniquely correct. If a child has only ever encountered the situation in a form where the correct action was supplied in advance, the appraisal step has never been practised. Whether closing this gap improves transfer is an empirical question that this paper does not answer. What this paper does is design an artefact in which that step is present.

UD — Understanding Duo — is an animated social skills curriculum for autistic children of roughly primary school age. The name refers to its two central characters, who appear together in every episode. The project

was originally designated ANIM-POD (Animated Interactive Modules for Personality Development in Children with Autism Spectrum Condition); we use UD throughout, since the artefact is the unit of interest rather than the programme under which it was developed. Its distinguishing property is structural: each episode is a decision tree. The narrative reaches a point where a character must act, the child is asked to choose among three options, and the episode continues along the branch corresponding to that choice. Branches reconverge; all of them resolve constructively. There is no losing path.

This paper reports the design of that artefact and the design knowledge abstracted from producing it. We follow design science research (DSR), which is well suited to work whose primary output is a constructed thing intended to solve a class of problem, and whose contribution to knowledge lies partly in the thing and partly in the principles that would let someone else build a comparable thing (Hevner et al., 2004; Peffers et al., 2007; Gregor & Hevner, 2013).

Three questions organise the paper:

What are the requirements for an animated social skills artefact that positions the child as chooser rather than viewer, given what is known about autistic learners?

What artefact satisfies those requirements, and how is it internally structured?

What transferable design principles does the artefact instantiate?

Section 2 reviews the relevant background. Section 3 sets out the research approach. Sections 4 and 5 report problem identification, solution objectives and the artefact itself. Section 6 abstracts the design principles. Section 7 demonstrates the design through a worked episode. Sections 8 and 9 discuss the contribution and its boundaries.

2. Background

2.1 Prevalence and the service gap

Autism is common enough, and unevenly enough served, that the design of scalable instructional material is not a marginal concern. A systematic review update covering studies published since 2012 reported a median global prevalence of 100 per 10,000, with wide variation both within and across regions (Zeidan et al., 2022). That variation is itself informative: much of it reflects differences in ascertainment and services rather than in underlying occurrence.

In India, where the present work is situated, prevalence in children has been synthesised by Chauhan et al. (2019). A country-wide scoping exercise conducted with 280 Indian respondents, followed by interviews with a stratified subset, found that community priorities diverge substantially from the research priorities that dominate high-income literature: self-help skills were ranked the highest training priority, on the reasoning that they underpin everything else (Dey et al., 2024). Material that a parent or a non-specialist teacher can run without extensive training therefore has a value beyond convenience.

2.2 Modelling-based approaches and the transfer problem

Video modelling has a favourable but thinner record than it is often credited with. Wang, Cui and Parrila (2011), synthesising single-case studies with hierarchical linear modelling, reported an average effect size of 1.22 for video modelling and 1.30 for peer-mediated approaches, with no significant difference between them, and found age to moderate effectiveness. The magnitudes are large, but the video-modelling estimate rests on five studies and carries a confidence interval from 0.65 to 1.78 — consistent with a substantial effect, but not precise enough to support strong claims. We cite it as encouraging rather than settled.

The qualification concerns what happens outside the training condition. Nejati et al. (2024), reviewing 52 controlled experiments through the FIELD framework (function, implement, ecology, level, durability), found transfer to be strongly influenced by age, autism severity, intervention tools, intensity, context and duration. Two of their findings bear directly on the present work. Cognitive interventions transferred better than behavioural ones, particularly at higher dose. And transfer was more pronounced for face-to-face interventions than for digital alternatives.

The second finding is inconvenient for a project that proposes a digital artefact, and we state it prominently rather than burying it. It does not, in our reading, argue against digital delivery as such; the reviewed digital interventions were overwhelmingly of the closed, response-specification kind. But it does establish a burden of proof, and it is one reason the adult support layer described in Section 5.6 is treated as integral rather than supplementary.

2.3 Social narratives

Social Stories, as developed by Gray (2010), describe a social situation and the responses appropriate to it, typically in short declarative sentences written from the child's perspective. The evidence is mixed in an interesting way. Qi et al. (2018), reviewing 22 single-case studies against What Works Clearinghouse standards, found that the intervention would not qualify as evidence-based on visual analysis but would qualify on non-overlap indices, and that effectiveness varied considerably between participants. Their reading — that social stories may work well for some children and not others — is more useful to a designer than a pooled average would be, because it directs attention to what varies.

The format's defining property is that it tells. A well-written social story leaves no ambiguity about what the child should do. That is precisely what makes it usable for children who find ambiguity aversive, and precisely what removes the appraisal step.

2.4 Animation, systemising and predictability

The case for animation is not simply that children like cartoons. The Transporters, developed at the Autism Research Centre, embedded real human facial expressions into animated vehicle characters on the reasoning that mechanical, rule-governed vehicles would engage the systemising preferences common in autism while the grafted faces carried the emotional content. Golan et al. (2010) reported improvements in emotion recognition across several levels of generalisation following a period of daily home viewing.

What this line of work establishes for our purposes is narrower than "animation works". It establishes that a controlled animated medium can carry socially relevant content in a form autistic children will attend to voluntarily and repeatedly, and that the controllability of the medium — no unexpected sensory events, no ambiguity in the presentation — is part of why.

Controllability matters because sensory hyper- and hypo-reactivity is a diagnostic feature of autism (American Psychiatric Association, 2013). An instructional medium that produces unpredictable loud or bright events is an instructional medium some children will refuse.

2.5 Interactive narrative and learner agency

Branching narrative is a mature form outside education and a much thinner one within it. Petousi et al. (2022) define interactive storytelling as a branching narrative in which users encounter different storylines and endings depending on choices made at decision points, and report a design for history education in which dialogic activity is deliberately located at those decision points. Their framing is useful: the decision point is not merely a navigation device but the place where the pedagogically interesting activity happens.

The motivational rationale is straightforward under self-determination theory, which holds that autonomy, competence and relatedness are basic psychological needs whose satisfaction supports self-motivation, and whose thwarting undermines it (Ryan & Deci, 2000). A choice that the learner experiences as genuinely theirs is autonomy-supportive in a way that a correct answer supplied in advance is not.

The risk, equally straightforward, is that choice implies the possibility of choosing badly. In a general-audience game this is the point. For an autistic child who may already have a history of social failure, an artefact that dramatises the child's wrong answer is not obviously therapeutic. This tension is what the non-punitive branching rule (Section 6, DP3) exists to resolve.

2.6 Participatory obligation

There is now a clear expectation in autism research that autistic people and those close to them shape the work rather than merely receive it. Fletcher-Watson et al. (2019), reporting a co-produced UK seminar series, identified respect, authenticity, assumptions, infrastructure and empathy as the topics on which a participatory community of practice depends, and noted candidly that meaningfully including autistic children, those with intellectual disability, and those who do not speak remains largely unachieved. We take this as a constraint on the design process, not an optional extra, and describe how it has and has not been met in Section 9.

2.7 The gap

Setting these strands together: modelling approaches have good within-condition evidence and a transfer problem; social narratives are closed by construction; animation is a proven carrier for autistic learners; branching narrative is a well-understood form that has barely been applied here; and choice architecture in this population raises a safety question that general interactive-narrative design does not have to answer.

No curriculum-scale artefact exists that combines these. That absence — a well-understood problem with no adequate existing solution — is what makes this a design problem rather than a study.

3. Research approach

3.1 Why design science research

DSR is appropriate when the primary output is a constructed artefact addressing a class of problem, and when the knowledge contribution consists partly in that artefact and partly in the abstracted principles behind it (Hevner et al., 2004). Hevner (2007) frames the work as three interlocking cycles: a relevance cycle connecting the artefact to its application environment, a rigour cycle connecting it to an existing knowledge base, and a central design cycle iterating between building and evaluating. That framing fits the present work: the application environment is autism education in a specific and resource-constrained setting; the knowledge base is the intervention and learning-design literature reviewed above; and the design cycle has run through several iterations of curriculum specification and screenplay drafting.

We use the six-activity process model of Peffers et al. (2007) — problem identification and motivation, definition of solution objectives, design and development, demonstration, evaluation, communication — to structure the account, noting that the model is nominal rather than strictly sequential and that our own process looped repeatedly between activities two and three.

This paper reports activities one to four, and six. Activity five, evaluation, is reported separately.

3.2 Scope of the artefact and its instantiation

A curriculum of 140 episodes cannot be animated within a doctoral programme. We therefore distinguish three nested scopes, and are explicit about which claims attach to which:

Scope A — the designed artefact. The full curriculum specification: 140 episodes across seven levels, each with skill objective, choice set, branch structure, complete screenplay, shot list, sound cue library and adult support notes. This is the design contribution.

Scope B — the instantiation. A subset of 14 episodes (two per level) produced as working animated software. This demonstrates that the specification is buildable and that the episode formula holds across skill types rather than only within one.

Scope C — the evaluated instance. A six-episode vertical slice drawn from three levels, used in empirical evaluation. [AUTHOR: this three-tier split is flagged in your project notes as needing formal locking; confirm the Scope B and Scope C counts before submission.]

Gregor and Hevner (2013) treat artefact abstraction levels explicitly as a matter for DSR authors to position rather than elide, and Gregor and Jones (2007) similarly distinguish the design theory from any particular expository instantiation of it. Selecting Scope B as a vertical slice across levels rather than a horizontal block within one level is a deliberate choice: it permits the claim that the episode formula generalises across skill domains, which a single-level block would not support. Two episodes per level is the minimum that allows within-level as well as across-level comparison while remaining producible at the animation quality the design requires.

3.3 Justificatory knowledge base

Following Gregor and Jones (2007), we state the justificatory knowledge on which the design draws, since design principles without stated rationale are difficult for others to evaluate or adapt.

Source	Contribution to design
Social learning theory (Bandura, 1977)	Modelling as an acquisition mechanism; vicarious consequence
Zone of proximal development (Vygotsky, 1978)	Graded scaffolding; the guide character as more-capable peer
Self-determination theory (Ryan & Deci, 2000)	Choice as autonomy support; non-controlling feedback
NDBI common features (Schreibman et al., 2015)	Shared control between child and adult; natural contingencies; developmentally sequenced targets
Structured teaching (Mesibov et al., 2005)	Visual structure, predictability, explicit sequence
Universal Design for Learning (CAST, 2024)	Multiple means of representation, action and expression, and engagement
Social narrative practice (Gray, 2010)	Situation framing; perspective sentences

Source	Contribution to design
Interactive storytelling (Petousi et al., 2022)	Decision points as the locus of pedagogical activity

4. Problem identification and solution objectives

4.1 Problem statement

Existing animated and narrative social skills materials for autistic children position the child as a recipient of a specified correct behaviour. This closes the situation appraisal step that real social competence requires, offers no autonomy support, and produces material that is structurally identical on every viewing, so that repeated exposure adds rehearsal but no new information.

4.2 Solution objectives

From the problem statement and the reviewed literature we derived six objectives. Each is stated with the source of the requirement, because objectives that cannot be traced to something are objectives that cannot be argued with.

SO1 — The child must act, not only watch. Every episode must contain at least one point at which the narrative cannot proceed without a decision from the child. Source: the appraisal gap (§2.2, §2.5); autonomy support (Ryan & Deci, 2000).

SO2 — No choice may be punished. All available options must lead to constructive, non-humiliating outcomes. Options may differ in how well they serve the character's goal, but not in whether the child is made to feel wrong. Source: the safety tension identified in §2.5; non-controlling feedback (Ryan & Deci, 2000); clinical caution regarding a population with elevated exposure to social failure.

SO3 — The medium must be sensorily predictable. No unannounced loud, bright or abrupt events. The child must be able to preview and control sensory intensity. Source: DSM-5 criterion B4 (American Psychiatric Association, 2013); structured teaching principles (Mesibov et al., 2005).

SO4 — Coverage must be curricular, not episodic. The artefact must span skill domains from foundational interaction through to safety, sequenced developmentally, rather than consisting of isolated units. Source: NDBI emphasis on developmental sequencing (Schreibman et al., 2015); practitioner requirement for a usable teaching order.

SO5 — Adults must be able to extend the episode. Every episode must carry guidance enabling a parent, teacher or therapist to reinforce the content in the child's ordinary environment. Source: the transfer problem, and specifically the finding that digital interventions have shown weaker ecological transfer than face-to-face ones (Nejati et al., 2024); NDBI implementation in natural settings (Schreibman et al., 2015).

SO6 — The artefact must instrument itself. Choice selection, replay, branch coverage and session behaviour must be recorded as a matter of design, not retrofitted. Source: evaluation requirements; the field's documented weakness in measuring what children actually do with an intervention.

5. The artefact

5.1 Curriculum architecture

UD comprises 140 episodes across seven levels of twenty episodes each. Levels are ordered by developmental prerequisite: each assumes the competencies of those preceding it.

Level	Domain	Representative content
1	Foundational social skills	Greetings, turn-taking, joining, requesting, ending an interaction
2	Emotional understanding	Recognising, naming and responding to emotional states in self and others
3	Daily living	Routines, transitions, self-care sequences, personal organisation
4	School	Classroom conventions, asking for help, group work, unstructured time
5	Community	Shops, transport, public spaces, encounters with unfamiliar adults
6	Friendship	Sustained peer relationships, disagreement and repair, reciprocity
7	Safety	Personal boundaries, unsafe requests, getting lost, seeking help

The sequence is not arbitrary. Emotional understanding (L2) precedes the applied levels because recognising a peer's state is a prerequisite for the friendship-level skills (L6) that depend on it. Safety sits last not because it is least important but because its episodes require the child to hold several prior competencies

simultaneously — reading an adult's behaviour, recognising discomfort, identifying a trusted person, initiating a request for help.

5.2 The episode formula

Every episode follows a fixed ten-beat structure. Fixity is deliberate: predictability of format frees the child's attention for the content, and is consistent with structured teaching practice (Mesibov et al., 2005).

Theme song. The same tune every episode, establishing the frame.

Arrival. A short, consistent beat in which the guide and protagonist enter the setting.

Today's mission. The guide frames the skill in a single line; a consistent mission-badge icon appears.

Observe. The social situation unfolds, including a held emotional or social cue beat in which the moment is allowed to sit rather than being narrated over.

Choose. The decision point. Three options, presented as image, audio and text, phrased as a direct question to the viewer.

See what happens. The chosen branch plays through to a realistic, non-punitive consequence.

Replay and compare. Optional: the child views an unchosen branch, followed by a simple side-by-side comparison.

Reflection. The guide asks a reflective question rather than stating a rule.

Try it today. A one-line real-world practice prompt, cueing the adult support material.

Goodbye song. A short closing jingle.

Beats 5 to 7 are the pedagogically active core; beats 1 to 4 and 8 to 10 are the invariant frame around it. The frame is not filler. Its purpose is to make every episode's shape known in advance, so that the only novel content the child has to process is the social situation itself.

Total runtime per episode is four to six minutes on any single path — longer than a conventional short-form module, because the invariant frame costs time that the design treats as well spent.

5.3 Decision-tree structure

The choice set is always three. Two is too few to make appraisal meaningful; four or more increases working memory load without adding pedagogical value at this age range. Options are presented in a fixed spatial arrangement across the curriculum so that position never carries meaning.

Two branch topologies are used.

Type A — equivalued. All three options are valid solutions to the situation. Consequences differ in texture rather than in quality. Level 1 Episode 1 (Saying Hello), where the options are Wave, Smile and Say hello, is of this type: each is a real greeting, and the episode's purpose is to establish that a social goal can be met in more than one way and that the child may choose the way that suits them.

Type B — differentiated. Options differ in how well they serve the character's goal. A queue episode might offer push forward, wait and leave the queue. Waiting serves the goal best. Crucially, the other two branches do not produce ridicule, punishment or catastrophe: pushing forward produces a mild, legible social consequence and a recovery, and leaving produces a different but liveable outcome. The child sees a difference in usefulness; the child is never shown having failed.

The distinction matters because a curriculum composed only of Type A episodes would teach that all social actions are equivalent, which is false, while a curriculum composed only of Type B episodes would make every decision a test. The curriculum specification assigns a type to each episode.

5.4 Characters

Two central characters carry the format. The protagonist is a young child of roughly five to six years, positioned as the figure with whom the viewer identifies — the one who encounters the situation and does not yet know what to do. The guide is an older child of roughly nine to ten years who accompanies the protagonist and addresses the viewer at the decision point.

The guide's register is the significant design decision. The guide asks rather than instructs. The line at the decision point is a genuine question, not a rhetorical one with a preferred answer signalled by tone. An older child rather than an adult was chosen because the near-peer relationship maps onto the zone of proximal development (Vygotsky, 1978) and avoids the compliance dynamic that an adult instructor introduces.

One further decision is worth stating because it was a revision rather than an initial choice. In an earlier design, the protagonist was a separate peer character whom the viewer watched, while a narrator framed the situation. That structure was collapsed so that the protagonist is the figure the viewer identifies with and also the figure who faces the dilemma and makes the choice. The effect is to move the design from peer modelling towards

something closer to self-modelling: the child is not observing someone else succeed, but occupying the position of the one who does not yet know what to do (Bandura, 1977). The recurring peer characters remain, but as the other person in the scene — the one the protagonist greets, shares with, or waits alongside — rather than as the model.
That recurring cast appears across levels. Recurrence is partly a production economy and partly pedagogical: familiar characters reduce the cognitive cost of each new episode, leaving more capacity for the social content.

5.5 Interaction mechanics

Three mechanics are embedded in the format rather than offered as features.
Replay and Compare. After a branch completes, the child may return to the decision point and select a different option. The interface presents this as exploration, not correction: the language is see what happens if rather than try again. This mechanic is the direct expression of the artefact's core claim, since it is the point at which a child can observe that a situation admits multiple outcomes.
Dual Perspective. In selected episodes the same moment can be replayed from another character's point of view, showing what that character saw and, where appropriate, what they were feeling. This addresses perspective-taking content directly rather than through narration.
Sensory Preview. Before an episode begins, the child is shown what sensory events it contains — whether there is a loud moment, a crowd, a sudden movement — and can adjust or mute them. This operationalises SO3 and, we would argue, is also a matter of respect: it treats the child's sensory response as information rather than as behaviour to be managed.

5.6 Adult support layer

The episode itself ends with a one-line real-world practice prompt (beat 9). Behind it, each episode carries a support sheet with three audiences addressed separately: parents, classroom teachers and support professionals. Each sheet states the skill objective, suggests two or three naturally occurring opportunities to prompt the skill in the child's day, gives language the adult can use, and specifies what the adult should not do — most commonly, should not require the child to perform the skill on demand.
This layer exists because the transfer evidence is pointed. Nejati et al. (2024) found ecological transfer — from the setting where training occurred to a different setting — to be the weakest of the five FIELD domains, with the large majority of reviewed studies showing none, and found digital delivery to fare worse than face-to-face. An artefact that does not reach into the child's ordinary environment has, by construction, ceded the variable that matters most. The adult layer is the artefact's answer to that, and it is the component whose usability most needs testing.

5.7 Instrumentation

Interaction logging is built into the episode player: which option was selected, in what order options were explored, whether Replay and Compare was used, session duration, episode completion, and any use of Sensory Preview controls. Logging is specified at the design stage rather than added later, both because retrofitted telemetry tends to be incomplete and because what to log is itself a design decision — the log schema encodes a theory of what counts as engagement.
Log data are pseudonymised at capture. Governance is described in the accompanying evaluation protocol.

6. Design principles

We abstract seven design principles from the design work. Each is expressed using the aim–context–mechanism schema of Gregor, Chandra Kruse and Seidel (2020), which requires a principle to state what it is for, in what circumstances it applies, and by what means it achieves its aim, together with the justificatory knowledge behind it. Principles stated only as slogans are not reusable; that is the schema's point.
DP1 — Principle of enforced appraisal Aim: To develop situation appraisal rather than response recall. Context: Instructional media for children who are being taught social behaviour. Mechanism: Structure each unit so that the narrative cannot proceed without an explicit decision by the learner from among plural valid actions. Rationale: Real social situations present options; instruction that supplies the answer never exercises the selection step (Bandura, 1977; §2.2).
DP2 — Principle of bounded choice Aim: To make choice meaningful without imposing load. Context: Learners in mid-childhood, including those with executive function differences. Mechanism: Present exactly three options, in a fixed spatial arrangement, in three modalities simultaneously (image, audio, text). Rationale: Multiple means of representation (CAST, 2024); minimisation of extraneous load; positional consistency prevents the interface from becoming a second problem to solve.

DP3 — Principle of non-punitive branching Aim: To preserve psychological safety while retaining consequence. Context: Learners with elevated exposure to social failure and to correction. Mechanism: Ensure every branch resolves constructively; differentiate options by usefulness towards the character's goal, never by the learner's worth; provide recovery within any less useful branch. Rationale: Controlling feedback undermines intrinsic motivation and autonomous regulation (Ryan & Deci, 2000); an artefact that dramatises a child's wrong answer risks teaching avoidance of the artefact. Note: This is the principle most likely to be contested. It trades some ecological fidelity — real social errors do sometimes carry real cost — for engagement and safety. We regard this as the correct trade at this age and in this medium, but it is a trade, and a fuller theoretical treatment of the trade-off, including the objections to it, is given in a companion paper [AUTHOR: cite companion conceptual paper once it has a DOI or preprint identifier].
DP4 — Principle of sensory pre-disclosure Aim: To prevent sensory-driven withdrawal from instruction. Context: Any digital instructional medium used by learners with atypical sensory reactivity. Mechanism: Disclose the sensory content of each unit before it begins and permit the learner to adjust it. Rationale: Sensory hyper- and hypo-reactivity is a diagnostic feature of autism (American Psychiatric Association, 2013); predictability of environment is central to structured teaching (Mesibov et al., 2005).
DP5 — Principle of the asking guide Aim: To position the learner as agent rather than as respondent. Context: Character-mediated instruction. Mechanism: Assign the decision-point address to a near-peer character who asks a genuine question, and prohibit prosodic or visual signalling of a preferred answer. Rationale: Zone of proximal development (Vygotsky, 1978); shared control as an NDBI common feature (Schreibman et al., 2015); autonomy support (Ryan & Deci, 2000).
DP6 — Principle of formula invariance Aim: To reduce the cost of encountering each new unit. Context: Curricula delivered across many units to learners who prefer predictable structure. Mechanism: Fix the beat structure, duration envelope, interface layout and interaction vocabulary across the entire curriculum; vary only content. Rationale: Structured teaching (Mesibov et al., 2005); reduction of extraneous processing so that capacity is available for the social content.
DP7 — Principle of designed instrumentation Aim: To make learner behaviour with the artefact observable. Context: Artefacts intended for empirical evaluation or iterative improvement. Mechanism: Specify the interaction log schema during design, deriving it from the constructs the artefact claims to affect. Rationale: Retrofitted logging captures what is technically convenient rather than what is theoretically relevant; DSR evaluation depends on observable artefact behaviour (Hevner, 2007).

7. Demonstration

Peffers et al. (2007) treat demonstration as showing that the artefact works to solve one or more instances of the problem, distinct from formal evaluation. We demonstrate through a structural walkthrough of a single episode.
Episode: Level 1, Episode 1 — Saying Hello Skill objective: Initiating a greeting with a familiar peer. Branch type: A (equivalued).
Theme song and arrival. The standard opening. The guide and protagonist walk into a neighbourhood park, midday, mild weather. The visual register is quiet: no crowd, no sudden movement, no abrupt sound.
Today's mission. The guide names the skill in one line and the mission badge appears.
Observe. A familiar peer is standing nearby, looking towards the protagonist. The protagonist stops. The moment is held — the situation is legible and the requirement to act is clear without being narrated over. This held beat is a production instruction as much as a narrative one; the temptation to fill the pause with explanation is exactly what the design resists.
Choose. The guide turns to camera and asks the viewer directly what the protagonist should do. Three options appear.

Option	Modality	Branch outcome
Wave	image, audio, text	The peer waves back; both smile; they walk on together
Smile	image, audio, text	The peer smiles in return; a brief shared moment; they walk on together
Say hello	image, audio, text	The peer says hello back; a short exchange follows

See what happens. The selected action animates through to a realistic, non-punitive consequence. Each outcome is warm and each is complete. The peer does not ignore, misread or reject any of the three.
Replay and compare. The child may return and select another option, then view the two outcomes side by side. Doing so is the pedagogical point of the episode: the child observes directly that three different actions all achieved the same social goal.

Reflection. The guide asks a reflective question rather than stating a rule, and does not rank the options — noting instead that different people prefer different ways of saying hello.
Try it today. A one-line practice prompt connects the skill to the child's own morning. The adult sheet suggests school gate arrival and the start of a therapy session as natural prompting opportunities, and advises against requiring a greeting on demand.
Goodbye song. The standard close.
The walkthrough demonstrates DP1 (the episode cannot proceed without a decision), DP2 (three options, three modalities), DP3 (all branches constructive), DP5 (the guide asks) and DP6 (the beat structure is the same one used by all 140 episodes).
One feature of this particular episode is worth noting because it is not repeated. Episode 1 is the only one in the curriculum that carries a name-reveal beat, in which the pairing of the two central characters is explained. Every subsequent episode uses the short standard arrival instead. The reveal is placed once and then left alone, on the reasoning that a format which re-explains itself is a format that has not been learned.

8. Discussion

8.1 Positioning the knowledge contribution

Gregor and Hevner (2013) classify DSR contributions along two dimensions — maturity of the problem domain and maturity of the solution domain — yielding four quadrants: routine design, improvement, exaptation and invention.
UD sits in the improvement quadrant. The problem is thoroughly understood: teaching social skills to autistic children in a way that transfers has been studied for decades, and the specific weaknesses of existing approaches are documented rather than speculative. The solution is not: no curriculum-scale artefact currently treats the episode as a navigable decision space with a non-punitive branch guarantee.
There is a secondary exaptation element. Branching narrative is a solved problem in entertainment and, increasingly, in general educational technology (Petousi et al., 2022). Bringing it into autism education requires more than transplanting it, because the population raises a constraint — DP3 — that general branching design does not have to satisfy, and that constraint materially changes the design. Where a general branching artefact uses consequence to teach, this one must teach while withholding negative consequence. That is a non-trivial redesign, not an application.

8.2 What the design principles offer beyond the artefact

The seven principles are stated at a level of abstraction intended to survive translation. DP3 in particular is not specific to animation or to social skills: it applies wherever choice-based instruction is offered to a learner group for whom the experience of being wrong carries disproportionate cost. That would include, plausibly, learners with anxiety-related presentations, learners with a history of educational failure, and adult learners in high-stakes retraining contexts.
DP4 and DP6 are similarly portable to any digital instructional medium used by sensory-sensitive learners.
DP1, DP2, DP5 and DP7 are more tightly bound to the present context and we make weaker claims for them.

8.3 Implications for practice

For practitioners, the artefact offers material that does not require the practitioner to be the source of the correct answer. A teacher can run an episode without knowing in advance what the child will choose, and the adult support layer gives them something to do afterwards that does not depend on their having specialist training.
For designers of similar systems, the principles are stated in reusable form and the artefact's failure modes are stated in Section 9.

8.4 What is not claimed

This paper reports a design. It does not claim that UD improves social skills, that branching narrative outperforms linear narrative, or that the design principles produce better outcomes than their alternatives. None of those claims can be made from design work alone, and we do not make them.
We claim that the artefact exists, that it is internally coherent, that it satisfies the stated solution objectives, that it is buildable, and that the design knowledge behind it is stated in a form others can use, adapt or contest.

9. Limitations

No empirical evaluation. The central limitation. The design is theoretically justified but not tested. Evaluation is reported separately and is the necessary next step.

The evidence on digital delivery is not favourable. As set out in §2.2, the most recent transfer synthesis found digital interventions transferring less well than face-to-face ones (Nejati et al., 2024). Our design response — the adult support layer — is a hypothesis about how to close that gap, not a demonstration that it closes.

Participation was consultative, not co-productive. Autistic children were not design partners in the sense Fletcher-Watson et al. (2019) advocate. [AUTHOR: state in one or two sentences exactly what input you obtained — practitioner consultation, caregiver input, autistic adult review — and its extent. Do not overstate.] This falls short of participatory best practice and we state it plainly rather than describing it in terms that imply more than occurred.

DP3 is a defensible trade, not a free good. Withholding negative consequence protects the child and may reduce ecological fidelity. Whether children who learn in a consequence-free environment behave differently when consequences are real is an open and important question.

Cultural specificity. The screenplays were written for a South Asian context and reflect particular assumptions about school routines, family structure, public space and greeting conventions. Some episodes will require substantial adaptation elsewhere, and the adaptation burden falls unevenly — Level 5 (Community) and Level 7 (Safety) are considerably more culture-bound than Level 2 (Emotional Understanding).

Heterogeneity is not addressed. The curriculum assumes a learner who can attend to a screen, process three simultaneous representational modalities, and make a selection. Autistic children who do not meet these assumptions — including many with co-occurring intellectual disability and many who are minimally speaking — are outside the artefact's current reach. This is a real exclusion and should not be softened.

Production feasibility. Only Scope B has been instantiated. Whether the specification holds across all 140 episodes at production scale is untested.

10. Conclusion

We have reported the design of UD, an animated social skills curriculum for autistic children in which episodes are structured as decision trees rather than linear narratives, and have abstracted seven design principles from that work. The central design commitment is that a child should be asked what to do and should be able to explore the consequences of more than one answer, without any answer being punished.

The contribution is design knowledge, positioned in the improvement quadrant of the DSR contribution framework. It is not an efficacy claim, and the paper should not be read as one. What we offer is an artefact that instantiates a different answer to a long-standing question about how social skills material should be structured, together with enough of the reasoning to let others build on it or argue with it.

References

1. American Psychiatric Association. (2013). Diagnostic and statistical manual of mental disorders (5th ed.). American Psychiatric Publishing.
2. Bandura, A. (1977). Social learning theory. Prentice-Hall.
3. CAST. (2024). Universal design for learning guidelines version 3.0 [graphic organizer]. CAST.
4. Chauhan, A., Sahu, J. K., Jaiswal, N., Kumar, K., Agarwal, A., Kaur, J., Singh, S., & Singh, M. (2019). Prevalence of autism spectrum disorder in Indian children: A systematic review and meta-analysis. Neurology India, 67(1), 100–104.
5. Dey, I., Chakrabarty, S., Nandi, R., Shekhar, R., Singhi, S., Nayar, S., Ram, J. R., Mukerji, S., & Chakrabarti, B. (2024). Autism community priorities in diverse low-resource settings: A country-wide scoping exercise in India. Autism, 28(1), 187–198. https://doi.org/10.1177/13623613231154067
6. Fletcher-Watson, S., Adams, J., Brook, K., Charman, T., Crane, L., Cusack, J., Leekam, S., Milton, D., Parr, J. R., & Pellicano, E. (2019). Making the future together: Shaping autism research through meaningful participation. Autism, 23(4), 943–953. https://doi.org/10.1177/1362361318786721
7. Golan, O., Ashwin, E., Granader, Y., McClintock, S., Day, K., Leggett, V., & Baron-Cohen, S. (2010). Enhancing emotion recognition in children with autism spectrum conditions: An intervention using animated vehicles with real emotional faces. Journal of Autism and Developmental Disorders, 40(3), 269–279.
8. Gray, C. (2010). The new Social Story book. Future Horizons.
9. Gregor, S., Chandra Kruse, L., & Seidel, S. (2020). Research perspectives: The anatomy of a design principle. Journal of the Association for Information Systems, 21(6), 1622–1652.

10. Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. MIS Quarterly, 37(2), 337–355.
11. Gregor, S., & Jones, D. (2007). The anatomy of a design theory. Journal of the Association for Information Systems, 8(5), 312–335.
12. Hevner, A. R. (2007). A three cycle view of design science research. Scandinavian Journal of Information Systems, 19(2), 87–92.
13. Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. MIS Quarterly, 28(1), 75–105.
14. Mesibov, G. B., Shea, V., & Schopler, E. (2005). The TEACCH approach to autism spectrum disorders. Springer.
15. Nejati, V., Peyvandi, A., Nazari, N., & Abadi, F. (2024). The effectiveness of social training in individuals with autism spectrum disorder (ASD): A systematic review and transfer analysis. Scientific Reports, 14, Article 32131. https://doi.org/10.1038/s41598-024-83953-9
16. Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. Journal of Management Information Systems, 24(3), 45–77.
17. Petousi, D., Katifori, A., Servi, K., Roussou, M., & Ioannidis, Y. (2022). History education done different: A collaborative interactive digital storytelling approach for remote learners. Frontiers in Education, 7, Article 942834. https://doi.org/10.3389/feduc.2022.942834
18. Qi, C. H., Barton, E. E., Collier, M., Lin, Y.-L., & Montoya, C. (2018). A systematic review of effects of social stories interventions for individuals with autism spectrum disorder. Focus on Autism and Other Developmental Disabilities, 33(1), 25–34.
19. Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. American Psychologist, 55(1), 68–78.
20. Schreibman, L., Dawson, G., Stahmer, A. C., Landa, R., Rogers, S. J., McGee, G. G., Kasari, C., Ingersoll, B., Kaiser, A. P., Bruinsma, Y., McNerney, E., Wetherby, A., & Halladay, A. (2015). Naturalistic developmental behavioral interventions: Empirically validated treatments for autism spectrum disorder. Journal of Autism and Developmental Disorders, 45(8), 2411–2428.
21. Vygotsky, L. S. (1978). Mind in society: The development of higher psychological processes. Harvard University Press.
22. Wang, S.-Y., Cui, Y., & Parrila, R. (2011). Examining the effectiveness of peer-mediated and video-modeling social skills interventions for children with autism spectrum disorders: A meta-analysis in single-case research using HLM. Research in Autism Spectrum Disorders, 5(1), 562–569. https://doi.org/10.1016/j.rasd.2010.06.023
23. Zeidan, J., Fombonne, E., Scorah, J., Ibrahim, A., Durkin, M. S., Saxena, S., Yusuf, A., Shih, A., & Elsabbagh, M. (2022). Global prevalence of autism: A systematic review update. Autism Research, 15(5), 778–790. https://doi.org/10.1002/aur.2696

Appendix A — Solution objectives mapped to design principles

Solution objective	Design principles	Artefact features
SO1 Child must act	DP1, DP5	Decision point; asking guide
SO2 No punished choice	DP3	Branch topology; reconvergence; Replay and Compare
SO3 Sensory predictability	DP4, DP6	Sensory Preview; fixed visual register
SO4 Curricular coverage	DP6	Seven levels; 140 episodes; fixed formula
SO5 Adult extension	—	Adult support layer (three audiences)
SO6 Self-instrumentation	DP7	Interaction log schema