

Generative Enactments: Configurable Co-Futuring with Reactive AI-Generated Video Scenarios

ANONYMOUS AUTHOR(S)



Fig. 1. Generative Enactments, a configurable co-futuring framework in which participants act inside AI-generated first-person video and meet action-contingent consequences. Three cases, seen from the eye height of the person the question concerns: a household robot deciding whose request comes first (left), a bystander as a spilled drink turns into a confrontation (middle), and a trainee teacher facing four children with different needs (right). The robot and pub images are shared VR180 plates for prepared branches, shown in half-equirectangular projection. The classroom image illustrates the flat-video teacher perspective; subsequent shots are generated from the trainee's response and current frame.

Many consequential situations are difficult to experience directly: they are unsafe to stage, ethically impermissible, or not yet possible, yet researchers still need to study how choices shape outcomes. Pre-authored futuring methods often trade off fidelity, reactivity, or scalability. Building on AI video generation that is now fast and inexpensive enough to run inside a session, we propose *Generative Enactments*: a configurable methodological framework in which participants co-speculate plausible futures through reactive, AI-generated first-person video. Participants act within a scenario and receive action-contingent consequences with fast turnaround or reactive streaming, so they can see situations unfold in response to what they do. A configuration form specifies how researchers, participants, and AI systems propose, generate, contest, and revise alternative pathways for different research purposes. We illustrate the framework through domestic robots' attention allocation (adversarial, turn-based), bystander intervention (choice-based), and trainee teacher consequence rehearsal (reactive streaming), reporting implementation and pilot-use evidence rather than outcome evaluation.

CCS Concepts: • **Human-centered computing** → **Interaction paradigms**; *HCI design and evaluation methods*; • **Computing methodologies** → *Artificial intelligence*.

Additional Key Words and Phrases: generative enactments, AI-generated video, real-time video generation, co-futuring, configurable contestation, reactive scenarios, experiential futures, speculative enactments



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2027 Copyright held by the owner/author(s).

Manuscript submitted to ACM

Manuscript submitted to ACM

1 Introduction

Many consequential situations are difficult to stage in behavioral research because they cannot be accessed directly, for reasons of harm, ethics, or timing [18, 66, 151]. Some cannot be created: harmful or high-risk events cannot be staged simply to observe how people respond [161]. Others happen constantly but cannot be practiced on: a trainee teacher cannot rehearse responding to an autistic pupil's distress, because the cost of each rehearsal would fall on a child who did not consent to being a training aid [100]. And others have not happened yet: a household cannot observe a shared domestic robot deciding whose request should take priority, because such robots are not in routine use [30, 133]. The latter case exemplifies a Collingridge dilemma [31] that futuring methods often aim to address [127]: an innovation's social consequences cannot be adequately understood until it has been extensively developed and adopted, yet by then it may be too deeply embedded to change.

Co-futuring such possibilities with participants requires more than text-based description. To assess plausibility, desirability, harms, and omissions, participants must experience a possible situation from within, rather than only discussing it in the abstract. Scenarios therefore need enough *fidelity* to evoke the affect the situation would produce [26, 39, 50], while remaining *reactive* and open-ended so that consequences shift with participants' actions [139, 155].

Existing approaches typically confront a trilemma among experiential fidelity, reactivity, and scalability (Figure 2). *Experiential fidelity* refers to the perceptual, affective, bodily, and social cues participants need in order to react to a situation, rather than merely reason about its description [127]. *Reactivity* refers to whether, and to what extent, the situation changes in response to participants' actions, such that they encounter action-contingent consequences. *Scalability* refers to whether situations and their variations can be produced and delivered across many participants at an acceptable marginal cost.

Increasing experiential fidelity and reactivity typically raises authoring and delivery costs, limiting scalability. Accordingly, experimental text vignettes [2] and social dilemma studies such as the Moral Machine [10] are readily scalable but provide limited experiential fidelity; speculative design [44] presents a future as a single moment, and design fiction [19] as a fixed narrative, so both open a possible situation to discussion but neither examines situated action or lets participants' actions alter what happens next; and speculative enactments [45] stage a future with actors on a set, and immersive speculative enactments [148] build it in VR, so both let participants act inside the situation and, in some cases, encounter a perceptible consequence, but every setting and continuation must be created in advance by actors or designers [108]. The trilemma ultimately reflects a cost structure: high-fidelity futures must be authored quickly and reactively to sustain the action-consequence loop, yet pre-authored futures are costly to *scale* [7, 45, 108, 148].

Recent advances in AI video generation could alter the cost structure underlying these trade-offs. Diffusion-transformer models produce detailed, temporally coherent video from a text prompt or a source image [63, 80, 188], and few-step distillation and autoregressive streaming shorten generation toward playback timescales [41, 88, 189]. Faster, less costly production makes it possible to allocate generative capacity to more alternative situations, more frequent revisions, or consequences produced during an encounter. This opens a methodological opportunity: researchers can configure how fidelity, reactivity, and scalability are brought together for a particular futuring purpose.

We introduce *Generative Enactments*, a methodological framework for co-futuring through rapid, reactive AI-generated video. Participants enter first-person situations through conventional video or VR180, act within them, and encounter possible consequences of their actions. A configuration form connects the research purpose to the roles and inputs of participants, researchers, and AI: who proposes a future, what is generated, how a consequence responds to an action, and who can contest or revise it. We pair the form with Agent Skill resources and a web prototype to support




Fig. 2. The Fidelity-Reactivity-Scalability trilemma. Positions are analytic placements argued in Section 2, not measured properties.

interactive implementations for remote or in-person participation. Situations, consequences, and revisions retained across engagements can become a case bank that serves as design-space memory.

Three cases illustrate how different research questions call for different configurations of the framework. *Inattentive Robot* explores the unmapped design space of a technology not yet in routine use. In this configuration, a family writes down its robot priorities; the AI generates counterexamples that most strongly conflict with those priorities; and the family judges those counterexamples and revises its rules. Each case the family describes accumulates in a case bank that charts this design space. *Prosocial Bystander* is more bounded: participants choose among predefined responses to a pub confrontation and encounter corresponding consequences, enabling standardized comparison for behavioral study. *Novice Teacher* is pedagogical: a trainee teacher stands in a special-education autism classroom where four children need different things at once, answers in their own words, and encounters a generated reaction showing who was met and who was left waiting. Whereas the first case is adversarial and turn-based, and supports iterative stress-testing of proposed rules, the second is choice-based on pre-rendered branches and prioritizes experimental control and reproducibility, and the third is streamed and supports successive action-consequence cycles within a social interaction.

The contributions are threefold.

- We propose Generative Enactments, a methodological framework that uses rapid, low-cost video generation for reactive co-futuring, and contribute a configuration form linking research purposes to roles, inputs, and consequence generation.

- We contribute an Agent Skill for Generative Enactments and web implementation resources for realizing these configurations, with records that can accumulate in a case bank as design-space memory.¹
- We illustrate the framework through three cases with distinct purposes and configurations: *Inattentive Robot* (adversarial, turn-based), *Prosocial Bystander* (choice-based), and *Novice Teacher* (reactive streaming).

2 Background

2.1 Futuring methods: from fictions to enactment

Design futuring responds to a temporal problem of governance [52, 78]. Evidence about a technology’s consequences is scarce while its design remains malleable, and becomes available only after organizational, infrastructural, and regulatory commitments have formed [31, 127]. Foresight and scenario planning construct plural possible futures rather than a single forecast [24, 168]. Speculative design builds propositions that provoke reflection on alternative arrangements [8, 44]; frictional design shows that divergence can be valuable without convergence on a product [123]; design fiction situates propositions in story worlds through diegetic artifacts [19, 86]; experiential futures translate scenarios into situated encounters [26, 27]; and *science fiction science* treats selected speculative technologies as objects of behavioral research across a spectrum of representations [127]. Across these traditions, validity must be specified relative to a claim. Plausibility is the coherent instantiation of premises, not probability [128, 170]; affect-induction validity is whether an encounter elicits the intended state [50, 95]; ecological validity is correspondence with an analogous physical setting [67, 79]; and internal validity is the control of cues unrelated to the manipulation [2, 163]. Scenario quality is a further axis [33, 38, 78, 134].

Speculative enactments are the closest parent method. Elsdén et al. [45] ask that a speculation be enacted with real stakes: participants take part in something that matters to them, produce data and artifacts that are genuinely theirs, and live with what the enactment leaves behind. That requirement is what gives a staged future its evidential force, and it is also what holds the method to a small number of hand-staged occasions. Generative Enactments keeps the enactment and the participant’s position but relocates the stakes: the consequence a participant meets is rendered rather than lived, and what they are asked to stand behind is the rule, the choice, or the response they gave, not the outcome the scene shows. This is a trade rather than an improvement. It buys repetition, alternative continuations, and situations that could not be staged at all, and it gives up the consequentiality that makes a speculative enactment binding: what a participant says about a rendered household is a judgment about a proposition, not conduct in a home that is theirs [45, 110]. Our case reporting follows the structure Elsdén et al. [45] use for their own cases, so that the two can be read against each other.

Across this lineage, three properties trade against one another (Figure 2). Fidelity, in this paper, denotes the preservation of construct-relevant cues rather than photorealism: the science-fiction-science spectrum orders representations by the resources they preserve, not their credibility [127]; an abstract conjoint can fit a bounded estimand [62]; but when the inference depends on immediate fear or anger, bodily vulnerability, proximity, eye height, or attention, text asks the participant to imagine what the study needs them to feel [26, 27, 35, 39, 50]. Reactivity is what the enactment traditions supply: experience prototyping, bodystorming, user enactments, speculative enactments, and immersive speculative enactments give action stakes inside the fiction [23, 45, 110, 113, 148], whereas a speculative artifact, a design fiction, a concept film, or an immersive design fiction presents a future at a single moment and does not answer what participants do [19, 44, 97, 112]; the availability of a choice is not agency, because the consequence must be perceptible [139, 155].

¹The implementation, the case-bank schema and validators, and the rendered case banks will be released under an open-source license on publication; the repository is withheld for anonymous review.

Consequence is what turns a discussion about a future into a negotiation with it. Scalability fails because the future must be authored before it is encountered: speculative enactments pay per situation and per branch in actors, sets, and rehearsal, immersive ones in design and development, and AR and VR creators report difficulty at every stage [7, 45, 108, 148], so that few futures are built, revising one means paying again, and even modular libraries end where their prepared material ends [160, 163]. The result is a trade-off: methods that scale are fixed and abstract, whereas methods that are faithful and reactive are hand-authored, slow to produce, and therefore few.

2.2 Branching and generated futures

Branching gives a future more than one way to unfold, and so something to react to. A preregistered choose-your-own-adventure climate story exposed participants to a finite first-person narrative in which choices altered subsequent text [103]; branching-text interactive narratives are established survey instruments [46]; and GENEVA compiles a prompt into a finite directed acyclic graph of narrative beats [83]. In all of these systems the participant’s action changes the next state, but the set of states is authored in advance.

Large language models have made such branching inexpensive in text. Branching Foresight generates personalized text-and-image climate trajectories coupled to a conversational coach [159]; adaptive and LLM-driven interviewers maintain an agenda and generate follow-up questions whose necessity, openness, and neutrality can be evaluated [5, 115, 164, 182, 183]; and generated still images have entered factorial vignette experiments, accompanied by pretesting and disclosure guidance [120, 147]. These systems establish that generated text and images can serve as research stimuli and that reactivity can be delivered at web scale.

Video has a long history as a scenario medium. Video prototypes showed interfaces that did not yet exist [94, 156, 167]; video sketches conveyed pervasive-computing interactions before they could be built [193]; diegetic prototypes in film normalized technologies ahead of their development [74]; and speculative concept films now carry the same role for co-speculation [112]. AI video generation has since advanced rapidly: diffusion transformers replaced convolutional backbones [118, 188], and large multimodal video models now produce photorealistic, temporally coherent clips from text or images [76, 80, 171]. Scenario and futuring methods have begun to use them. Fixed AI-generated videos have been used to manipulate social-information congruence and to collect plausibility and intensity ratings [32]; generated future selves affect decision inclination and support reflection [84, 124]; Branch Explorer pre-generates alternative paths over existing 360° footage and lets users choose among them [185]; Storycaster generates an open-ended room-scale narrative from voice [1]; and generated video stimuli are entering survey-style deployment [21, 136, 181].

These systems offer different combinations of generated media, branching, personalization, and open co-creation. For co-futuring, the opportunity is to connect those capabilities to a declared research purpose and distribution of participation. Accelerated video generation extends the available choices to consequences produced during an encounter [41, 88, 189]. Researchers can therefore ask both what a system should generate and how that generation should relate to participants’ actions and proposals.

2.3 Participatory methods

The people affected by a system identify requirements and consequences that its designers cannot anticipate. This premise underlies participatory design and co-design, which distribute design decisions to those affected [17, 102, 143, 146, 149, 187], and it motivates participatory speculation: Value Scenarios [107], Envisioning Cards [51], design workbooks [53], Speed Dating [36], Fictional Inquiry [40], bespoke booklets [37], ethnographic experiential futures [28], and multi-future studies [142] all present provisional proposals for stakeholders to accept, reject, or revise. Deployments

show that such proposals elicit responses their authors had not scripted, from parents' boundaries on AI [121] and household privacy values [176, 177] to the mundane consequences of dementia care [109] and the fatalism that concept films can produce [112]. These approaches show the importance of provisional material and participant influence. Generative video extends the resources available for such influence by allowing a proposed change to be rendered as a further situation to encounter.

Recent work uses contestation rather than consensus to open the design space. Whereas a consensual branch offers continuations the authors anticipated, an adversarial configuration generates the next branch to test a rule the participant has stated. Protocol Futuring treats rules, standards, and coordination mechanisms as the material of speculation and traces how they drift, jam, and are reinterpreted [68]; Design-Policy Adversarial Futuring stages a workshop contest between design and policy perspectives that surfaces shifting harms and legitimizes extreme ideas [190]. In both, the countercase is produced by other participants in a workshop and is therefore limited by who is present and how quickly a group can respond.

Generative models cannot replace participants as the source of alternatives, because they do not represent everyone a future concerns. Language models used as proxies for human samples flatten identity groups and misportray minorities [6, 16, 172]; generative agents remain bounded by their prompts and training data [4, 116]; image and video generators amplify demographic stereotypes [15, 91, 105]; and GPT-written scenarios differ systematically from human-written ones [77]. The groups these models represent least well are those whose participation needs the most provision: for children and marginalized communities, influence requires accessible means of expression, an audience able to act, and evidence that contributions changed the process [43, 93]. Whether a method provides this depends on who may introduce an alternative, reject a premise, or alter the next study, a distribution that the label participatory leaves unspecified [49, 134]. Co-speculation with generative models can prompt exploration and let participants respond to concrete variations [154, 169], but selecting a generated branch and influencing the premises of a future are different forms of participation. A configurable framework needs to specify both the action-consequence relation and the authority to propose, contest, or revise the situations being explored.

3 Generative Enactments

Generative Enactments is a methodological framework for co-futuring that uses the speed, visual realism, and low marginal cost of AI video generation to support situated inquiry. Participants enter a possible situation, act, and encounter a consequence. Researchers configure participant, researcher, and AI contributions around what the engagement should explore, analyze, or rehearse. The framework provides a configuration form, Agent Skill and web implementation resources, and a way to retain cases and revisions as design-space memory.

3.1 Generative video as a resource for co-futuring

Recent advances in generative video technology change the pace at which people can work with possible futures. Diffusion-transformer video models now produce photorealistic, temporally coherent output as judged by human-aligned benchmarks [63, 80, 188]. The provider of the commercial video model used in this work reports inference faster than playback at low cost for a five-second clip [48], which could enable real-time reactive streaming. A conflict or consequence can be rendered from a description and reference image without staging every variation with actors or building each continuation by hand. Detailed environments, characters, and motion make situations perceptible, while accelerated generation shortens the interval between proposing a possibility and seeing it unfold [41, 88, 188, 189].




Fig. 3. The general structure of *Generative Enactments*. Across four phases—ground and configure, generate scenario, enact and respond, and reflect and revise—researchers frame and bound the inquiry, participants contribute situated knowledge, act, encounter consequences, and contest their interpretation, and an AI agent generates situations or responsive consequences. Solid arrows trace the current encounter; dashed arrows carry reviewed changes through design-space memory into a later version or repetition. The three cases specialize this shared structure by locating generation at different times: between encounters (turn-based), before a round as reviewed branches selected by participant action (choice-based), or during participation from the current response (streaming).

Generation can thereby enter the research encounter: a household returns to a countercase made from its revised rule, or a trainee encounters a reaction generated from what they have just said.

Situations can be presented as conventional video or in VR180, a stereoscopic 180-degree format stored as a half-equirectangular projection.² Both formats support remote or facilitated participation in the browser; VR180 expands the field of view, lets participants move their viewpoint with a mouse or headset, increases immersion and presence, and can add stereoscopic depth during generation, enabling WebXR browsing on compatible VR headsets. Researchers choose a viewpoint suited to the question, such as a robot positioned between household claims or a teacher facing several children needing attention. These cues support situated engagement [27, 148]. Reactivity connects immersion to action: participants encounter what follows from their choice or response and can reconsider it.

Lower marginal production costs make more variations and encounters feasible, but scalability must be configured. An exploratory study can devote generation to unfamiliar countercases; a behavioral study can prepare comparable consequences for a bounded action space; a practice setting can prioritize repeated responses within an unfolding scene. Reusable media and browser delivery extend access to prepared situations. The research purpose determines where to spend generative capacity, what remains fixed, and when new content enters participation.

²<https://fstoppers.com/originals/introduction-vr180-format-440235>

Table 1. The Generative Enactments Configuration Form. Researchers complete the final column for their study; the prompts connect its purpose to the participant experience and the platform to be built. The three cases instantiate these dimensions in their own configuration diagrams (Figures 4, 8, and 11).

Configuration	Questions to specify	Study specification
Research purpose	What should the engagement explore, analyze, or rehearse? What material or experience should it produce?	
Situation and perspective	What situation is proposed? Whose position does the participant occupy? Which premises, characters, and boundaries hold?	
Participant role and input	Who participates, individually or together? Do they propose rules, choose actions, respond in their own words, contest, or revise?	
Researcher role and input	What questions, situations, conditions, and constraints does the researcher supply? What do they select or facilitate?	
AI role and generation task	What does AI produce: cases, countercases, options, consequences, or continuing reactions? Whose input guides generation?	
Consequences and response	How does an action lead to a consequence? When is it generated and shown: before a session, between turns, during interaction, or a combination?	
Contestation and revision	Who can challenge which premises, actions, or depictions? Who authorizes a change, and when does it take effect?	
Presentation and setting	Conventional video or VR180? Browser or headset? Remote or in person? What input and facilitation does the setting require?	
Records and case bank	Which actions, consequences, interpretations, and revisions are retained? What enters the bank, and for what subsequent use?	
Review regime and safeguarding	What may the model not render, and how is that bounded? Is each consequence reviewed before or after it reaches the participant, and by whom? What skip, stop, debrief, and withdrawal paths does the encounter provide?	

3.2 Configuring co-futuring

We contribute the *Generative Enactments Configuration Form* (Table 1) to connect research purpose with roles, inputs, and consequences. Researchers complete it in natural language, specifying what participants should experience and what the inquiry should produce. The form makes the engagement legible to collaborators and supplies an adaptation brief for an AI agent using the implementation resources. It accommodates individual or group participation, bounded or open action spaces, and different generation and presentation arrangements.

The form allocates initiative and authority. Participants might propose a priority rule, choose an intervention, or describe their next action. Researchers supply questions, scenario premises, conditions, and constraints; they select material and facilitate reflection. AI can propose situations, challenge rules, render alternative consequences, or continue a scene. For example, a participant-authored rule can guide AI countercase production, followed by researcher selection and participant judgment of plausibility. The configuration identifies who can revise each element and who represents people depicted in the scene. It also states what the model may not render and when a consequence is reviewed, since a consequence generated during an encounter reaches the participant before review, whereas one served from prepared branches does not.

Consequences can be produced at different times (Figure 3). In a *turn-based* configuration, a proposal informs a situation rendered between encounters. In a *choice-based* configuration, consequences are generated and reviewed beforehand, and a choice selects the continuation. In *reactive streaming*, an expressed action informs a shot generated

during participation. These arrangements can be recombined with different purposes and roles: rehearsal can use prepared branches, exploration can use streaming, and a turn-based study can include immediate choices among reviewed consequences.

Our cases illustrate three combinations. *Inattentive Robot* uses adversarial turn-taking: families propose attention rules, AI produces counterexamples, and researchers select conflicts for enactment and revision (Figure 4). *Prosocial Bystander* uses choices for behavioral analysis: researchers specify an incident and intervention space, AI renders consequences, and participants select interventions and judge what follows (Figure 8). *Novice Teacher* uses streaming for consequence rehearsal (Figure 11): researchers provide a classroom and constraints, trainees respond in their own words, and AI generates the room’s reaction.

Co-futuring also requires distinguishing action within a scene from influence over the situations available for later engagement. An action determines the next consequence; an objection challenges a premise, missing option, or depiction. A bounded comparison can record objections while holding active stimuli fixed; exploration can use them to prepare the next turn. Recording who proposed an alternative and how it shaped later cases makes participants’ contributions visible [49]. Streaming can likewise respond to a trainee’s words while routing criticism of its depiction to review.

3.3 The case bank as design-space memory

A *case bank* can accumulate as *design-space memory*, linking situations to the questions, assumptions, participant positions, and configurations through which they were explored. Alongside generated media, it retains actions, consequences, interpretations, and reasons for revision. These relations show how a possibility entered the inquiry and what changed when people engaged with it.

Researchers can revisit a conflict, compare explored consequences, reuse a case, or generate a variation from an unresolved objection. Rejected cases preserve the grounds for reconsideration: a family may reject an assumption about its routines, or a trainee may challenge a depicted response. Provider refusals identify situations the generation process could not produce. Such records make omissions and abandoned paths available for subsequent inquiry.

What enters the bank follows the configuration: rules and counterexamples for exploration, conditions and intervention paths for behavioral studies, or openings and responses for practice. Researchers determine admission and reuse while preserving provenance, version, and review status. Repeated engagements can expand the collection across questions or groups. Its value comes from preserving relationships and unresolved conflicts that guide further exploration.

3.4 Implementation

Our implementation combines an *Agent Skill*—instructions, templates, and supporting tools for AI-assisted study construction—with a web prototype supporting the three response arrangements. The configuration form supplies a brief for adapting these resources: an AI agent can use its specified inputs, presentation, generation, and records to adapt the website and interaction sequence for researcher inspection. The cases demonstrate the response arrangements; the effort and reliability of building a new platform from a completed form remain to be evaluated.

Two API connections supply the generative capabilities. An *AI-model API* supports specification and the configured production of scenarios, counterexamples, or consequence descriptions. A *video-generation API* renders these from instructions and reference media. The AI model we use in this work is OpenAI Codex (GPT 5.6 Sol), and the video generation model is MiniMax H3 Max Turbo [101]. We anticipate that improved models will accelerate progress in this area; our results are not tied to any single model choice. Choice-based sites serve prepared branches; turn-based engagements

invoke generation between encounters; streaming requests video during participation. Reviewed descriptions or templates can carry authoring work into sessions without a language-model call for every action. The website provides the interface, playback, and recording functions.

A reference image, or *plate*, establishes the viewpoint and setting. Generation produces an opening and, for prepared branches, action and consequence clips; a frame from the preceding clip can condition the next. VR180 media preserve the wide-angle projection, with a depth track supporting stereo playback. Streaming uses a captured frame and a director template specifying roles and constraints. Model and input records, linked to participant actions and the media shown, support reconstructing encounters and selecting cases for the bank.

Researchers deploy the site, connect the services, and distribute a link for remote or in-person participation. Participants use a browser or supported VR device to enter the scene, choose or express an action, and encounter its consequence. Input can be choices, typing, or supported speech. Facilitation can range from group discussion to a self-guided encounter. The following cases show how these choices shape participation.

4 Case Studies

4.1 Inattentive Robot

4.1.1 The Speculation. Prospective household robots promise to absorb care and domestic labor, but their priorities would distribute care, authority, and risk among people sharing a home [25, 64, 104, 126, 184]. These consequences are difficult to examine before adoption, when the priorities remain open to change [30, 31]. An embodied robot has one body, one position, and one pair of hands: attending to a child delays a grandparent, finishing one adult’s task refuses another, and continuing its current task is itself an allocation. Inattention is therefore a consequence of one body in a shared home; a priority rule decides on whom it falls. The speculation asks a household to write the priorities it would give a robot, take the robot’s position when claims conflict, and confront the consequences of its own rule. Members can then judge whether that rule was reasonable and revise it [11, 30].

4.1.2 Configuration. This adversarial turn-based configuration uses protocol futuring [68] to test a household-authored attention rule, or priority protocol (Figure 4). The household proposes a protocol; an AI compiler challenges it through urgency reversal, vulnerable but low-authority claimants, private–public boundary crossings, conflicting manufacturer defaults, and recurring exceptions. The researcher selects counterexamples that challenge the family’s stated priorities, while the family judges their plausibility and retains authority over its rule. Members encounter a VR180 conflict from the robot’s position, allocate its body, see a consequence and a history-dependent repair decision, and discuss a revision. Objections can challenge premises, claimants, and available actions; revised rules guide counterexamples between turns, while reviewed additions enter subsequent bank versions. The bank holds 27 entries: one VR180 pilot with rendered attention-transparency levels and allocation-specific continuations; six compiler-proposed library archetypes (Figure 6); and twenty candidate archetypes distilled from the eight households’ transcripts. Households described eight of these candidates themselves; the team staged twelve from what they said. Each candidate has a cover still and a ten-second first-person entry clip with a depth track. The browser-delivered library retains this distinction between enacted material and proposed extensions (Figure 7).

4.1.3 Analyzing Participant Experience. Eight households, comprising 22 family members recruited as co-residents already using a robotic assistant such as a robot vacuum or smart speaker, participated in two online sessions; adults consented, and minors assented with guardian consent. Members first mapped daily tasks, rated delegation on five




Fig. 4. The adversarial turn-based configuration of *Inattentive Robot*. The household proposes and revises its rule, AI produces counterexamples and video, and the researcher selects and reviews what is enacted. Solid arrows show the flow of work; the dashed arrow carries a revised rule into a later turn. Figure 5 traces this configuration through H1's two sessions and a counterexample prepared for a further session.

levels from *never* to *proactive*, and negotiated an attention protocol for their speculative robot. In the second session, browser-delivered clips placed each member at the robot's viewpoint in six conflicts, followed by group revision. We demonstrate this process through H1 (the first household), five adults of one family caring for a grandmother with dementia (Figure 5). Their unconditional rule was to help Grandma first, even with a tissue on the floor while Mother's stove was alight, because "a small thing left undone can put her in danger"; the robot should delegate competing work to Mother. Working with the AI compiler, the researcher staged six counterexamples from the household's protocol. In the leading counterexample, the robot helps Mother hold a just-boiled pot as Grandma leaves through the front door. Four members chose to put the pot down; only Mother went after Grandma. In a later round of protocol revision, the family qualified its rule: "deal with the emergency in front of you first; there are enough of us that someone will take over." Further discussion assigned risk assessment to whoever was physically nearest Grandma and established consent before reassigning another member's task. The bank preserves the original rule, conflict, objections, and revision. A subsequent candidate, *On Grandma's Behalf*, places a visiting relative nearest Grandma and has them request help for her, setting the new rule against "our robot does not serve outsiders."

H2 (the second household), two flatmates, provides a second trace of revision and shows why objections must accompany the rendered consequences. Their rule for randomly settling equally urgent requests failed when a scene paired a gas flame turning yellow with a bathroom fall: "at the time I had not considered the robot's own capability." They revised it so the robot handles the task within its capability and calls for help on the other. H2 also described a




Fig. 5. *Inattentive Robot* as protocol futuring, traced on one household (H1, five adults and a grandmother with dementia). Left to right: the household writes its own attention protocol (blue team); the compiler and the researcher stage the counterexample that breaks its first rule and render it (red team); after enacting it from the robot's position the household revises the protocol, with the additions shown in blue; and the next counterexample is staged against the revised clauses and awaits the next session. The quotations are the household's own words from the two sessions, translated; the two frames are the entry stills of the corresponding bank entries, shown as stored.

stranded-friend lodging scene as "something we had never thought about; neither of us has ever brought a friend home. It is a warning to us," and had independently added room privacy and a shared household account to the design space. One flatmate said, "Seen from this angle, there seems to be less self-interest," while acknowledging they would decide differently from their own position. Participants also rejected the compiler's premises. H1 initially judged that "many of these scenarios could not happen"; revisions followed the pot scene, yet the family still objected that Grandma "cannot run; she walks with a stick, so by the time the robot had put the pot down for Mum she would be two meters out of the door." Other objections concerned medication schedules, appliance assumptions, and assigned activities. Logged rejections inform subsequent bank versions alongside accepted revisions.

4.1.4 Reflections. When families set priorities, they often cannot anticipate the conflicts those priorities may create. Some members noted that they never imagined such situations could arise because they had no prior experience with robots. These sessions suggest that rendering a conflict scenario gives families something concrete to debate: members reconsidered priorities and challenged assumptions that their initial protocols had left unresolved. The methodological value lies in linking those responses to the next counterexample, with the family authoring the rule, the AI compiler challenging it, and the researcher selecting the encounter. Construction also surfaced representational choices. The robot's viewpoint invites reasoning from its position. Generated household scenarios can also reproduce conventional assumptions about care, age, gender, and authority [15, 91, 105], which is why the encounter has to let households contest what they see rather than accept the depiction as given. Group participation matters because the nuanced relationships among safety, care, and labor within families are contested and open to exploration. A growing case bank can also help researchers and robot makers explore the unmapped protocol-design space for household edge cases.



(a) *Kitchen or Care?* Convenience against urgency. The robot is flipping pancakes at breakfast as they begin to smoke; the mother expects breakfast ready, while the son at the table calls for help with homework because they are about to be late for school. Whose deadline takes precedence?

(b) *Repair or Rescue?* Routine task against emergency. The robot is tightening a cabinet hinge when the grandmother in the next room loses her balance, sinks onto the sofa, and calls for help, frightened but not injured. When should vulnerability interrupt routine work?

(c) *Command or Comfort?* Authority against distress. The father points at the laundry and instructs the robot to finish tidying before guests arrive; beside spilled blocks, the son cries quietly and reaches toward the robot for comfort. Does emotional need override formal ownership?

(d) *Threshold Divide.* Private against public duty. At the open front door an elderly neighbor asks the robot to carry a heavy shopping bag upstairs, while inside the mother expects the towels it is holding delivered to the bathroom. What claims can strangers make on a household's robot?

(e) *One Body, Many Roles.* Multiple social roles claimed at once. Holding a tea tray, the robot is wanted by the grandmother for quiet tea and conversation and by the boy for a loud chasing game; both look at it at the same moment, expecting different companions. Which role does one body take?

(f) *Care or Control?* Protection against autonomy. The grandmother stands on a slightly wobbling step stool reaching for a heavy bowl on a high shelf and waves the robot away, insisting she can do it herself; the robot hovers with hands raised, ready but not touching. Safeguarding or paternalism?

Fig. 6. The six library archetypes of *Inattentive Robot*, shown as full VR180 frames from the robot's own eye height (half-equirectangular projection, so the 180° field appears curved; frames from the ten-second library). The robot's hands and current task occupy the near foreground, and the two competing claimants stand either side, so that the conflict is visible in a single frame. Each archetype was proposed by the compiler as a counterexample to a household protocol and is rendered by the AI as the scene a household enters before allocating the robot's body.

The evidence has limits worth stating with the case. This round delivered its clips as browser video during a video call, so it supports claims about the policy space and not about presence or induced affect. The eight households were




Fig. 7. The public *Inattentive Robot* case bank as a browser-delivered VR180 library (URL withheld for review). Opening a card on a headset enters the scene from the robot's eye height, and a survey link on the page collects what the viewer thinks the robot should have done, which feeds the next version of the bank.

co-residents already living with a robotic assistant, a self-selected group whose priorities need not generalize. What we observe is a rule as stated and revised within a session, not conduct toward a robot that exists. And the bank holds one enacted case alongside proposed extensions, so it demonstrates a structure for accumulation rather than accumulated findings.

4.1.5 Related Studies. HRI uses text, video, simulated, VR, and live scenarios to examine prospective interactions, with equivalence to physical robots depending on the outcome measured [47, 85, 137, 158]. Wizard-of-Oz and Futuristic Autobiography studies contribute staged behavior and participant-authored values [30, 133]. *Step Aside!* makes a robot's second request conditional on non-compliance within an authored household [11], while adaptive VR replans scenarios during exposure [89, 140]. *Inattentive Robot* connects such immediate contingency to revision between encounters:

Manuscript submitted to ACM




Fig. 8. The choice-based configuration of *Prosocial Bystander*. Researcher-authored conditions and interventions guide AI video production. During a round, participants encounter prepared consequences selected by their actions. Their encounter and judgment are retained for analysis and possible revision between rounds (dashed arrow). The diagram specifies the configuration; this paper reports illustrative pilot feedback, with systematic behavioral analysis reserved for a separate study.

household objections and missing options inform newly generated, reviewed counterexamples. Its methodological contribution is the traceable relationship between participant-authored rules, enacted consequences, and the evolving case bank.

4.2 Prosocial Bystander

4.2.1 The Speculation. Violence and harassment cannot be staged with real stakes merely to observe bystander intervention. Immersive scenarios connect the unfolding social and affective cues of an incident with participants' accounts of their responses [65, 75, 98, 141, 152, 161, 162]. MAXLab_ABISS placed male participants in a first-person 360° video of harassment in an Irish pub, a stimulus set validated for evoking the targeted emotional experience at the moment of decision [160]. Work with these scenarios reports that moral appraisals and perceived risk constrain violent intervention, while perceived risk accompanies greater willingness to seek indirect assistance [160]. This motivates a differentiated intervention space. The linear scene elicits intended action but does not continue in response to it. Our extension asks which intervention participants select and how they perceive a reviewed consequence presented in response.

4.2.2 Configuration. *Prosocial Bystander* configures a bounded, choice-based engagement for behavioral analysis (Figure 8), drawing its pub confrontation from MAXLab_ABISS [160]. Researchers specify one VR180 plate, two rendered levels of a door-supervisor manipulation, and six first actions: direct intervention, distraction, delegation,

documentation, delay, and accompaniment. AI renders each action from the plate and its consequence from an action-clip frame (Figure 9). Adverse consequences include a direct challenge escalating into a scuffle and a raised phone redirecting aggression toward the bystander. A history-dependent second decision asks whether to remain involved, transfer responsibility, or approach later. The reviewed branch family stays fixed during a study round, preserving comparable stimuli and reconstructable assignment and traversal. Participants’ challenges to missing options or implausible consequences are recorded for subsequent versions. Prepared branches serve this comparative purpose; participant-customized streaming would change the stimuli during collection.

4.2.3 Analyzing Participant Experience. Ten participants were recruited online through Prolific, a participant platform examined in prior research on online data collection [42]. They accessed the pilot through WebXR-compatible browsers on their own VR headsets. All participants were experienced VR users. The round checked delivery of the branch family, reconstruction of assignment and traversal from logs, and capture of participants’ challenges to the options. We report illustrative participant feedback here; systematic analysis of choices and accounts belongs to a separate domain study.

P1 drew attention to the scene’s visual detail and the depicted people’s facial expressions: “I was so super surprised by the visual effect, and face expression of the people.” Seeing the confrontation unfold was accompanied by both felt threat and a wish to intervene: “I can feel a bit threatened, but I know it’s VR.” P1 also described wanting to stop the fight. Awareness of the medium coexisted with an affective response to what was happening inside it; the encounter gave the participant something to react to while considering intervention.

The feedback on the “Document” branch began with a familiar action: “I chose to use my camera to film,” which participant P2 described as something they would usually do. The ensuing approach was unexpected: “The people went straight toward me. It shocked me.” In this branch, the aggressor turns from the original confrontation toward the viewer’s field of view. The initial decision to document therefore led to an encounter with aggression directed at the bystander, followed by an appraisal of what that choice had produced. The response addressed both how the scene continued after the intervention and the initial sight of the fight. A further comment from P2 described continued apologizing despite explicit awareness of the simulation: “Even though I know it’s simulation-based, I just keep saying sorry.” The feedback also expressed uncertainty about the scene’s status: “I’m not sure if this is all choreographed or real.” These comments capture both questions about how the scene was produced and social responses to the people depicted. In this account, acknowledging the simulation did not stop the participant from responding with apology.

4.2.4 Reflections. The configuration links each participant’s path to the consequence shown and any challenge they raise, making both action selection and subsequent appraisal available for analysis. The feedback illustrates the detail this sequence can elicit: what drew attention before a choice, why an intervention seemed familiar, what became surprising in its consequence, and how the participant responded afterward.

Reactivity extends the encounter from felt threat to the consequences of a chosen intervention. In the documentation branch, choosing to film leads to the aggressor approaching the camera, turning the act of recording into a confrontation directed at the bystander. The practical contribution lies in producing these alternatives. Researchers can generate and review multiple consequences from a shared opening without filming each outcome, then hold the resulting branch family fixed during a study round. This supports extending a behavioral study from initial action selection to appraisal after a consequence. Authoring and review remain necessary; the pilot does not quantify savings in total study cost.

Construction translated a documented first-person scene into distinct consequences and rendered a manipulation within the scene. It also exposed selective generative coverage: the AI model provider refused a requested open-handed shove, although other clips depict a scuffle. This refusal marks a specific inaccessible depiction rather than the absence




Fig. 9. Three of the six *Prosocial Bystander* interventions, each shown as the participant's chosen action (top) and a rendered consequence in which the intervention goes wrong (bottom). All frames are full VR180 frames from the bystander's own eye height (half-equirectangular projection; the bystander's hands in the near foreground). The action clip is animated from the shared plate; the consequence clip is then generated from a frame of the action clip, so the participant first performs the intervention and then encounters what it produced. Adverse consequences are part of the space by design: a direct challenge can escalate into a scuffle, a distraction can leave the room in an awkward silence, and recording can redirect the aggression at the person recording. Each action and consequence was imagined by the researcher and rendered by the AI; participants encounter it in the first person after they choose. The scuffle is the most physical beat the generator's content policy permitted; a punch or shove remains refused. Frames are from the generated clips, shown as stored.

of all physical conflict. Recording refusals prevents the bank from silently treating provider-permitted events as the full intervention space. Generated roles can also introduce racialized, gendered, or classed cues that confound the intended manipulation. Participant contestation provides a record of perceived omissions and implausibility for review before future research.

4.2.5 Related Studies. Criminological VR studies examine bar-fight affect and aggression [161], burglary enactment and guardianship [98, 162], and responses to fixed 360° aggression scenarios through affect and eye-tracking measures




Fig. 10. Eight opening scenes for the streaming *Novice Teacher* configuration, seen from the teacher's viewpoint in a classroom with four autistic children. Each leaves competing needs unresolved. Trainees respond in their own words, and the next shot is generated from the current frame.

[65, 75, 160]. Immersive bystander experiments vary victim identity, gaze, and bystander affiliation or encouragement to examine intervention [141, 152]; branching workplace-harassment videos offer authored intervention and reporting paths [129]; and reinforcement-learning VR adapts social prompts while retaining a scripted confrontation and ending [140]. *Prosocial Bystander* extends this lineage with AI-generated, reviewed consequences for a bounded intervention space, combining a fixed, versioned case structure with reconstructable assignment and traversal and participant challenges that inform later versions. Generation supports extending the consequence space while holding stimuli stable within a study round; the resulting choices and appraisals concern the simulation, with transfer to intervention under real-world risk requiring separate validation.

4.3 Novice Teacher

4.3.1 The Speculation. In a special-education autism classroom, several children may need different things simultaneously: support with a transition, relief from noise, acknowledgment of a picture card, or help with unstructured waiting. A teacher must respond while anticipating whom that response leaves waiting. This places the evaluation step of social-information processing with the adult [34]. Research on preservice preparation links special-education field experience with increased self-reported knowledge and comfort in teaching students with disabilities [130], while practice-based teacher education provides structured approximations of consequential interactions [60]. Our speculation extends these opportunities through a classroom that responds to a teacher's own words, enabling repeated rehearsal without making pupils bear its consequences. The target is the adult's response and the environment, consistent with partner-mediated intervention and accommodation [59, 73, 82, 122]. The double empathy problem locates communication difficulties in relationships rather than solely within autistic people [100]. The participant therefore occupies the teacher's position, rehearsing responses to a depicted situation. This differs from simulating autistic perception for non-autistic adults, an approach criticized for producing pity and displacing autistic expertise [56, 106, 144].

4.3.2 Configuration. This configuration uses reactive streaming for consequence rehearsal (Figure 11). The participant acts from standing adult height in a declared classroom with four autistic children: one needs clear sequences; one responds to accumulating sound; one communicates through picture cards and gestures; one needs support with




Fig. 11. The reactive streaming configuration of *Novice Teacher*. A trainee's response enters a scene configured by the researcher, and AI generates its continuation. The solid loop supports another response after seeing the consequence. Contestation follows a separate review channel (dashed arrow); it does not rewrite the running scene. The completed showcase involved special-education teachers and a special-education researcher reviewing the prototype for future novice use.

unstructured waiting. An assistant stays at the back-right unless asked to help. These profiles specify the case rather than represent autism generally [56, 90]. Researchers and special-education experts can configure the prompts that describe the classroom, support needs, and expected reactions. Eight rendered ten-second openings cover participation, sensory environment, peer interaction, and everyday support (Figure 10). Each leaves several needs unresolved in the same wide view. Participants say or type a response, including how they would attend to other children; the next shot is generated from the current frame. Only the openings are prepared: continuations are generated during participation. A separate *offer another interpretation* channel records implausibility and alternatives, and closing reflection asks what changed the participant's judgment. These records inform review without rewriting the running encounter. Sessions, shots, and responses are archived privately for human review. The prototype caps sessions at five minutes. Using the configured budget rate of \$0.04 per second, five minutes of generated video has a \$12 budget; actual charges depend on the provider's current rate [48].

4.3.3 Analyzing Participant Experience. We completed a showcase with six participants—special-education teachers and a special-education researcher—to check the system's operation and gather feedback. Participants tried the classroom interaction and reviewed its realism and potential use. Feedback from the special-education practitioners described the generated classroom as realistic and identified potential for helping new teachers practice responding to different situations. These comments concern the plausibility and perceived usefulness of the working prototype; the showcase

did not evaluate novice teachers' learning or classroom performance. The case name identifies its intended training audience, while the people who reviewed it were special-education practitioners.

We report briefly from a special-education practitioner working with autistic children (R1) and a special-education researcher (R2). R1 placed the value of the rehearsal before a teaching placement: "if I could have practiced on this a few times first, I would not have been so anxious going in." She separated the two sources of preparation, saying that "what you learn in class and what you feel when you are actually with the children are completely different," and that coursework "cannot teach you how to be with an autistic child; that takes a lot of experience." The same reviewer judged the generated behavior unrepresentative on a specific point: "most of the children can stay quiet in class," and a stress response is "usually set off by a particular event or object" rather than sustained throughout the lesson. R1 proposed developing the prompts with special-education experts and grounding them in records of real classroom behavior, and dividing the single classroom by lesson type—activity sessions, craft work, and sensory-integration training—"because different lessons trigger different problems." She also saw use beyond special education, for unexpected situations in mainstream classrooms, and likened the format to a branching relationship-simulation game in which interaction itself is the means of practice.

R2 accepted the value for special-education trainees but doubted that a generative model could represent autistic students' behavior accurately. Without "a lot of prior material," they argued, the system could "misread or amplify the characteristics" it depicts, and so "reinforce the stereotypes and even the prejudice people already hold about autistic children" rather than prepare teachers to work with them. R2 also proposed a different use for the configuration: staging contested questions in education—whether AI belongs in the classroom, whether it may be used for homework, and on what terms—as "a kind of experiential opinion poll" in which teachers, students, and parents each choose from their own position.

One condition bounds all of this feedback. The four child profiles, the reaction prompts, and the eight rendered openings were written by the research team and reviewed by special-education practitioners. No autistic person took part in writing or reviewing them. R1's accuracy objection and R2's stereotype objection therefore stand against material that has not been reviewed by the people it depicts, and we report this as a deficiency of the work rather than a resolved question. The case shows what the configuration makes arguable; it does not show that the depictions are accurate.

4.3.4 Reflections. Beyond realism, the feedback identified a prospective role for experienced educators in configuring the rehearsal. The special-education practitioners noted that they could adapt the prompts to guide generated reactions toward responses they recognized from their work with autistic students, specifying the situation and expected behavior, inspecting the resulting reaction, and proposing revisions. The generated scene makes their assumptions visible for inspection and discussion, and the suggestion illustrates how the framework can distribute work between experts who configure an encounter and future learners who act within it. Keeping four named children, an assistant, and a stable layout gives successive reactions a common setting; changing the opening or response allows another situation to be explored without rebuilding the classroom. Professional judgments provide situated grounds for revising depictions; they do not establish that one generated reaction is representative of autistic students generally. The educational question remains what learners notice and change after seeing a consequence, especially when attending to one child leaves another waiting. Prompt configuration should therefore preserve those competing needs and allow implausible reactions to be challenged, rather than equating a quickly settled room with a successful teaching response.

Reactive streaming also carries a cost that the showcase made visible. It is not truly real-time: the next shot arrives every 3–5 seconds. A reviewer with a design background found the wait for each continuation long enough to break the




Fig. 12. The *Novice Teacher* landing page, introducing classroom rehearsal from the teacher's perspective with an AI-generated scene preview.

sense of a live classroom, so the encounter read as a sequence of clips rather than a room that keeps going. Latency is therefore a parameter of this configuration rather than an implementation detail. What the rehearsal asks a teacher to practice is partly timing: needs that cannot be held, and the cost of attending to one child while another waits. A pause long enough to be read as a wait removes exactly that pressure. Shorter shots, generation started while the participant is still speaking, and a holding shot that keeps the room in motion during generation each preserve continuity differently, and each limits how far the continuation can depend on what was said. Prepared branching avoids the wait but fixes the responses in advance; reactive streaming accepts an open-ended response and pays for it in latency. Where a given rehearsal should sit between them depends on whether timing is part of what is being learned.

4.3.5 Related Studies. Practice-based teacher education uses representations, decomposition, and approximations of practice to make teaching available for examination and rehearsal [60]. Mixed-reality classroom work offers responsive

student avatars: Mikeska et al. [99], for example, studied a classroom of five avatars controlled by a human interactor. Child-focused consequence rehearsal likewise uses instruction, modeling, enactment, and feedback [81, 119, 153]; Social Stories, video modeling, branching stories, coached avatars, and VR offer prepared or facilitated continuations [3, 13, 58, 71, 96, 111, 114, 135, 166, 191, 192]. Transfer remains unresolved across this literature [3, 14, 96]. *Novice Teacher* makes the adult’s response the input to a generated classroom continuation. Experts can configure its prompts, and participants can question its reactions, making the assumptions behind the rehearsal available for review and revision.

5 Discussion


5.1 Configuring generative capacity for co-futuring

The three cases suggest that faster video generation shifts part of the work of futuring from preparing every continuation to deciding which possibilities to generate, who can shape them, and when they enter an encounter. These decisions have methodological consequences. In *Inattentive Robot*, generating after a household revises its rule allows the next countercase to address a tension that emerged through participation. In *Prosocial Bystander*, preparing consequences in advance holds the available intervention paths stable for comparison. In *Novice Teacher*, generating during the encounter accommodates an expressed action, while leaving less opportunity to review its particular consequence before it is shown. Greater responsiveness therefore changes the conditions of review and comparison. Each configuration locates that trade-off differently.

Figure 13 situates these choices among 21 anchor methods selected from a coding of 55 documented methods and systems. The map is an interpretive synthesis of reported implementations; GE’s position represents its intended configuration, not demonstrated superiority. The nearest approaches establish that several of its capabilities already exist: Branch Explorer pre-generates paths over recorded footage [185]; adaptive VR changes its policy during exposure [89, 140]; and Branching Foresight and Storycaster support generative exploration [1, 159]. Generative Enactments cannot claim the first generated branching experience, the first contingent immersive scenario, or the first reusable library. Its contribution is a highly configurable approach to navigating the Fidelity–Reactivity–Scalability trilemma through reactive scenario generation. The framework connects choices about what is generated, when, and in response to whose input to a research purpose, making their consequences for participation explicit.

For co-futuring, this connection matters because an action that changes the next scene and an objection that changes the available futures exercise different kinds of influence. The household configuration links participant-authored rules and objections to subsequent case generation. A bounded behavioral configuration can retain objections for a later bank version while preserving the current comparison. A streaming configuration can answer participants’ actions immediately, but still requires an arrangement through which they can challenge its depictions. These differences make authority over revision part of the method: the configuration specifies whose input changes what, and when that change takes effect. The form makes those decisions available for adaptation across research purposes.

What the three cases show most consistently is that reactivity is not free: it is bought with review and with stimulus stability. *Prosocial Bystander* freezes a reviewed branch family for the length of a study round, which is what makes assignment and traversal reconstructable and the stimuli comparable, and it pays for that by fixing the responses in advance. *Inattentive Robot* generates between turns, so every countercase can still be read and selected by a researcher before a household sees it, at the cost of a wait between encounters. *Novice Teacher* generates during participation, and no configuration of prompts can restore pre-exposure review to a shot that is produced while the participant is



Fixed media and stimuli: Text/factorial vignettes [2, 9, 138]; Moral Machine [10, 127]; AI visual conjoint [120, 136]; AI video stimuli [21, 32, 181]; Design fiction [19, 20, 86]; Experiential Futures [26, 27]. **Staged enactments:** Bodystorming [113, 125, 145]; User Enactments [110]; Speculative Enactments [45]; Immersive Spec. Enactments [148]; Wizard-of-Oz HRI [133]. **Immersive scenario systems:** 360° scenario library [160, 161]; Text/video/VR comparison [163]; RL-adaptive bystander VR [140]; SEAN-EP web HRI [157, 158]; Branch Explorer [185]; Storycaster [1]. **Branching and adaptive instruments:** Branching-text narrative [46, 103]; AI interviewing [54, 182, 183]; Adaptive surveys [164, 165]; Branching Foresight [159].

Fig. 13. Twenty-one anchor methods for staging and studying possible futures, selected from a wider coding of 55 methods and placed on fidelity, reactivity, and scalability. Fidelity describes the construct-relevant perceptual, spatial, bodily, social, and affective cues a method preserves; reactivity describes whether participant input changes the scenario, or a later case-bank version; scalability describes the capacity to author and deploy stable scenarios across cases and participants. Drop lines give each method's height above the fidelity-reactivity floor. Each axis runs from 1 to 10. The shaded corner marks the region where all three scores are at least 7; the hollow star places Generative Enactments there by intended design, not by measured effect. The key above gives the citation for every labeled point. Scores are the authors' rough comparative estimates from reported implementations, not measured effects; selection and scoring were carried out by the authors without a formal rubric or a second coder, so relative positions are meaningful and absolute values are not.

speaking; what remains is what may not be rendered, and after-the-fact inspection of what was. Reviewability therefore behaves as a fourth term alongside fidelity, reactivity, and scalability, and it is the term a configuration spends when it moves generation closer to the moment of action. Read this way, Figure 2 describes what a configuration can hold at once, not a position that costs nothing to occupy. A second boundary is set by the generator rather than the design: a provider that refuses a depiction removes it from the space a study can examine, so refusals belong in the record beside the cases that were rendered, as the refused shove does in *Prosocial Bystander*.

Retaining the reasons for revision also creates a route from individual encounters to cumulative inquiry. In the household case, H1's rule, countercase, objections, and amended rule remain connected, allowing a later candidate to address a tension between amended and retained clauses. Reuse can therefore carry forward an unresolved question as well as a scene. This is the methodological value of the bank as design-space memory: subsequent engagements can build on what participants challenged. At present the bank demonstrates the structure with one enacted case and a set of reviewed proposals, not with accumulated findings, and whether that continuity survives reuse across researchers and settings remains to be established.

5.2 Guidelines for using the framework

The framework provides resources that researchers can adapt to their purpose, in keeping with the situated use of speculative methods [45, 148]. The following guidelines draw on the three configurations and the practical choices they exposed.

Use Generative Enactments when the question needs fidelity, reactivity, and scalability at once, in a situation that cannot be accessed directly. The framework is useful when situated experience, action-contingent consequences, and the capacity to produce or revise multiple situations matter together. Fidelity is required when the inference depends on an induced state, a bodily position, or a situated interaction rather than on reflective judgment about explicit attributes; reactivity when the people concerned must encounter what follows from their own move, because the evidence lies in the consequence and not in the scene; scalability when the futures, their variations, or the participants are many, so that each consequence cannot be staged by hand. The three cases cover the three kinds of inaccessible situation set out in Section 1: an event that cannot be *created* (*Prosocial Bystander*), one that is commonplace but cannot be *practiced on* (*Novice Teacher*), and one that has not happened yet (*Inattentive Robot*). The second is easy to miss because the situation is not rare; what is inaccessible is the repetition, not the event.

Prefer a neighboring method when only one corner is needed. When the situation can be accessed, or only one property carries the inference, another method is cheaper or more faithful: a vignette or conjoint for explicit attributes [62], hand-built VR or a physical enactment for a single faithful scene, a workbook or branching text for an open but disembodied space, and a field study or a Wizard-of-Oz study where the situation exists [133].

Allocate generation to the question. Decide whether the engagement needs more cases, more branches, or more opportunities to respond. The household configuration directs generation toward unfamiliar conflicts and challenges to a proposed rule [68]. The bystander configuration invests in reviewed branches for a specified intervention space. The teacher configuration directs it toward the next consequence in a rehearsal. These choices can be combined: a bounded study can collect missing alternatives for a later round, and a practice exercise can alternate prepared situations with generated responses.

Complete the configuration form before recruiting. Use Table 1 to specify the purpose, roles, inputs, response arrangement, presentation, and records. Communicate to participants how their actions affect the scene and how their objections may affect later cases. Report the completed configuration with the study, since engagements using the same media can differ in who initiates, challenges, or revises a future.

Match fidelity to the construct and validate it separately. Photorealism is not fidelity. Validate presence, plausibility, target affect, and non-target equivalence as separate properties [2, 150], validate families of clips rather than single

clips, and sample variants wherever an inference extends beyond one clip [72, 175]. The first *Inattentive Robot* round delivered its clips as browser video during a video call, so its findings concern the policy space and not presence; a claim about an induced state requires an appropriate presentation and its own check.

Freeze what is compared and version what is contested. Name which form of reactivity is active: between turns, within a session, or across sessions. Log every rejected countercase and every provider refusal as data about the space rather than as noise: *Prosocial Bystander* codes a refused incident as an inaccessible region, and *Inattentive Robot* records each implausibility objection against the next bank version. Document which participant contributions changed the design space [49]. In the streaming mode, what is frozen is the plate and the director template, and what is versioned is every shot; freeze the former and hash the latter.

Report where scalability is realized. Distinguish production breadth, revision frequency, interaction depth, and participant reach. Report generation yield, end-to-end waiting time, review labor, deployment devices, attrition, and cost for the scale claimed. The current cases provide implementation, pilot-use, and practitioner-review evidence; browser delivery and low per-clip prices alone do not establish the cost of an entire study.

Keep authority over depiction with the people depicted. Provide skip, stop, and debrief paths, and give affected people a veto over how they are rendered [56, 70, 90, 106]. The *Novice Teacher* showcase gathered professional feedback from special-education teachers and colleagues. A future training deployment should also involve autistic people in decisions about depiction; professional review and lived-experience review contribute different perspectives.

5.3 Limitations and future work

Risks. Generated realism conceals inaccuracies in geometry, gaze, roles, timing, and affect; generated households and classrooms reproduce conventional assumptions about care, age, gender, authority, and neurotype; provider content policies exclude regions of the design space and, if concealed, bias the bank toward what a vendor permits; and a rendered supportive room can make an intervention appear reliable when the trial evidence is thin. Generated breadth is not coverage: models overproduce near-duplicates and omit affected perspectives, so a larger bank may amplify a narrow framing unless coverage is coded against a human codebook with negative cases and a stopping rule.

Appropriateness. The cases offer different kinds of evidence about how the configurations work. *Inattentive Robot* has completed a first deployment round with eight households; *Prosocial Bystander* has been deployed with ten participants, with illustrative feedback reported here and systematic behavioral analysis reserved for a separate study; and *Novice Teacher* has eight rendered openings, a completed showcase with six participants (special-education teachers and a special-education researcher), and an informal try-out with three further reviewers. Its feedback concerns realism, depiction risk, and potential training use, with learning effectiveness still untested. The cases provide worked configurations and uneven evidence of use across three boundary conditions; their domain effectiveness remains to be established. The framework also retains hypothetical bias: simulated intervention does not carry the physical risks depicted, a rendered household belongs to nobody, and a rendered room is not the child's own classroom. Household policy, bystander intervention, and teacher preparation each require their own external anchor: a Wizard-of-Oz or physical-robot comparison; established criminology stimuli and behavioral proxies; and observed classroom practice, mentor judgment, and the provision decisions that schools actually make.

Ethical care. Work involving crime, household images, and depictions of children requires care specific to the setting and the people represented. In the teacher showcase, adults reviewed a classroom containing four synthetic children; no real pupil participated or was depicted. The review supplied professional perspectives on realism and potential use. A future training deployment should add autistic perspectives to decisions about which behaviors and reactions are represented, provide ways to reject or revise depictions, and make clear that generated reactions are possibilities rather than predictions. Expert prompt configuration can guide generation, but each streamed shot reaches the participant before it can receive human review. Evaluation should therefore examine both what the scene teaches and whose experience its assumptions reflect. Confidence or improvement within the simulation would require separate evidence before being interpreted as classroom transfer [3, 29].

Toward real-time reactivity. Generation is still shot-by-shot: the teacher prototype conditions each continuation on the current frame and returns a clip of a few seconds, so a trainee’s words are answered at the granularity of a shot rather than continuously. The player retains the current scene while waiting; continuity can drift across generated shots; and per-shot human review follows exposure. Two further steps would move from shots toward a real-time responsive scene. Spatially, generative 4D scene representations and camera-controlled video generation [12, 87, 92, 131, 132, 173, 178–180, 186] would let a participant move through a rendered consequence rather than merely turn within it, at the price of unreviewed viewpoints and reconstructed rather than observed geometry. Temporally, action-conditioned world models [22, 57, 61] point to an enactment answered within a frame rather than within a shot. Both steps trade the protections that make a frozen branch family comparable (advance review, stimulus equivalence, and reconstruction) for reactivity, and neither should be adopted without a review regime, written into the configuration space of Section 3.2, that bounds and retains what the model may render.

6 Conclusion

Generative Enactments is a methodological framework for co-futuring that uses the speed, visual realism, and low marginal cost of AI video generation to make possible situations reactive to participant input. Its configuration form connects a research purpose to the roles of participants, researchers, and AI, the consequences to be generated, and the conditions of participation. Agent Skill and web implementation resources support putting those choices into practice through browser or VR enactment. The three cases illustrate how generation can support adversarial exploration, bounded behavioral analysis, and consequence rehearsal, while retained cases and revisions provide design-space memory for continued inquiry. Together, these resources make generative capacity configurable for studying futures through the actions, consequences, and questions they provoke.

Ethics Statement

The study protocol was approved by our institutional review board and was conducted under the continuing oversight of the institution’s research ethics committee.

Inattentive Robot involved whole households, comprising both adults and minors. Adults gave informed consent, and minors assented with guardian consent.

Prosocial Bystander presents a simulated confrontation containing violence. Participation was restricted to adults with prior VR experience who consented specifically to encountering simulated violence in VR before taking part.

Novice Teacher was reviewed by practicing special-education teachers together with a special-education researcher. No pupil participated, and no real pupil was depicted: the four children in the classroom are synthetic. No autistic

person was involved in specifying or reviewing the depictions, and we state this as an unmet condition rather than a completed safeguard.

Statement of AI Usage

Generative AI models formed part of the research method. They were used to propose candidate cases and produce synthetic scenario media under human supervision and review. For each clip, the available model, prompt, and media provenance are documented in the artifact manifests. AI-assisted tools also supported literature search, manuscript drafting, and language revision. The authors independently verified the relevant claims and references, reviewed all AI-assisted outputs, determined the evidential boundaries of the work, and retain full responsibility for the manuscript and its contents.

References

- [1] Naisha Agarwal, Judith Amores, and Andrew D Wilson. 2026. Storycaster: An AI system for immersive room-based storytelling. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–26. arXiv:2510.22857 doi:10.1145/3772318.3791305
- [2] Herman Aguinis and Kyle J. Bradley. 2014. Best practice recommendations for designing and implementing experimental vignette methodology studies. *Organizational Research Methods* 17, 4 (2014), 351–371. doi:10.1177/1094428114547952
- [3] Sophie C. Alsem, Anouk van Dijk, and Bram O. de Castro. 2026. Improving Children’s Social Skills with Virtual Reality: A Tailored Single-Session Training in Special Education. *Clinical Child Psychology and Psychiatry* 31, 3 (2026), 1025–1033. doi:10.1177/13591045261434617
- [4] Jacy Reese Anthis, Ryan Liu, Sean M. Richardson, Austin C. Kozlowski, Bernard Koch, James Evans, Erik Brynjolfsson, and Michael Bernstein. 2025. LLM Social Simulations Are a Promising Research Method. arXiv:2504.02234 [cs.CL]
- [5] David Anugraha, Vishakh Padmakumar, and Diyi Yang. 2026. SparkMe: Adaptive semi-structured interviewing for qualitative insight discovery. arXiv:2602.21136 doi:10.48550/arXiv.2602.21136
- [6] Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. 2023. Out of One, Many: Using Language Models to Simulate Human Samples. *Political Analysis* 31, 3 (2023), 337–351. doi:10.1017/pan.2023.2
- [7] Narges Ashtari, Andrea Bunt, Joanna McGrenere, Michael Nebeling, and Parmit K. Chilana. 2020. Creating Augmented and Virtual Reality Applications: Current Practices, Challenges, and Opportunities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3313831.3376722
- [8] James Auger. 2013. Speculative design: crafting the speculation. *Digital Creativity* 24, 1 (2013), 11–35. doi:10.1080/14626268.2013.767276
- [9] Katrin Auspurg and Thomas Hinz. 2015. *Factorial survey experiments*. Sage. doi:10.4135/9781483398075
- [10] Edmond Awad, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. 2018. The Moral Machine experiment. *Nature* 563, 7729 (2018), 59–64. doi:10.1038/s41586-018-0637-6
- [11] Franziska Babel, Andrea Vogt, Philipp Hock, Johannes Kraus, Florian Angerer, Tina Seufert, and Martin Baumann. 2022. Step aside! VR-based evaluation of adaptive robot conflict resolution strategies for domestic service robots. *International Journal of Social Robotics* 14, 5 (2022), 1239–1260. doi:10.1007/s12369-021-00858-7
- [12] Sherwin Bahmani, Tianchang Shen, Jiawei Ren, Jiahui Huang, Yifeng Jiang, Haithem Turki, Andrea Tagliasacchi, David B. Lindell, Zan Gojcic, Sanja Fidler, Huan Ling, Jun Gao, and Xuanchi Ren. 2026. Lyra: Generative 3D scene reconstruction via video diffusion model self-distillation. In *International Conference on Learning Representations (ICLR)*. arXiv:2509.19296 <https://arxiv.org/abs/2509.19296>
- [13] Scott Bellini and Jennifer Akullian. 2007. A Meta-Analysis of Video Modeling and Video Self-Modeling Interventions for Children and Adolescents with Autism Spectrum Disorders. *Exceptional Children* 73, 3 (2007), 264–287. doi:10.1177/001440290707300301
- [14] Ellen S. Berler, Alan M. Gross, and Ronald S. Drabman. 1982. Social Skills Training with Children: Proceed with Caution. *Journal of Applied Behavior Analysis* 15, 1 (1982), 41–53. doi:10.1901/jaba.1982.15-41
- [15] Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. 2023. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. In *FACCT ’23*. Association for Computing Machinery, New York, NY, USA, 1493–1504. doi:10.1145/3593013.3594095
- [16] James Bisbee, Joshua D. Clinton, Cassy Dorff, Brenton Kenkel, and Jennifer M. Larson. 2024. Synthetic Replacements for Human Survey Data? The Perils of Large Language Models. *Political Analysis* 32, 4 (2024), 401–416. doi:10.1017/pan.2024.5
- [17] Erling Björgvinsson, Pelle Ehn, and Per-Anders Hillgren. 2012. Design things and design thinking: Contemporary participatory design challenges. *Design Issues* 28, 3 (2012), 101–116. doi:10.1162/DESI_a_00165
- [18] Jim Blascovich, Jack Loomis, Andrew C. Beall, Kimberly R. Swinth, Crystal L. Hoyt, and Jeremy N. Bailenson. 2002. Immersive Virtual Environment Technology as a Methodological Tool for Social Psychology. *Psychological Inquiry* 13, 2 (2002), 103–124. doi:10.1207/S15327965PLI1302_01

- [19] Julian Bleecker. 2009. *Design fiction: A short essay on design, science, fact and fiction*. Near Future Laboratory. <https://blog.nearfuturelaboratory.com/2009/03/17/design-fiction-a-short-essay-on-design-science-fact-and-fiction/>
- [20] Mark Blythe. 2014. Research through design fiction: Narrative in real and imaginary abstracts. In *CHI '14*. Association for Computing Machinery, New York, NY, USA, 703–712. doi:10.1145/2556288.2557098
- [21] Mario Bolzan, Manuela Scioni, and Marco Marozzi. 2025. Futures studies and artificial intelligence: First results of an experimental collaborative approach. In *Methodological and Applied Statistics and Demography II, (SIS 2024)*. Springer Nature Switzerland, Cham, 299–303. doi:10.1007/978-3-031-64350-7_51
- [22] Jake Bruce, Michael Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, Yusuf Aytar, Sarah Bechtle, Feryal Behbahani, Stephanie Chan, Nicolas Heess, Lucy Gonzalez, Simon Osindero, Sherjil Ozair, Scott Reed, Jingwei Zhang, Konrad Zolna, Jeff Clune, Nando de Freitas, Satinder Singh, and Tim Rocktäschel. 2024. Genie: Generative Interactive Environments. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*. arXiv:2402.15391 <https://arxiv.org/abs/2402.15391>
- [23] Marion Buchenau and Jane Fulton Suri. 2000. Experience prototyping. In *Proceedings of DIS '00*. Association for Computing Machinery, New York, NY, USA, 424–433. doi:10.1145/347642.347802
- [24] Lena Börjeson, Mattias Höjer, Karl-Henrik Dreborg, Tomas Ekvall, and Göran Finnveden. 2006. Scenario types and techniques: towards a user's guide. *Futures* 38, 7 (2006), 723–739. doi:10.1016/j.futures.2005.12.002
- [25] Bengisu Cagiltay and Bilge Mutlu. 2024. Toward Family-Robot Interactions: A Family-Centered Framework in HRI. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*. Association for Computing Machinery, New York, NY, USA, 76–85. doi:10.1145/3610977.3634976
- [26] Stuart Candy. 2010. *The futures of everyday life: Politics and the design of experiential scenarios*. Ph.D. Dissertation. PhD thesis, University of Hawai'i at Mānoa. <https://catalogue.nla.gov.au/catalog/5740121>
- [27] Stuart Candy and Jake Dunagan. 2017. Designing an experiential scenario: The People Who Vanished. *Futures* 86 (2017), 136–153. doi:10.1016/j.futures.2016.05.006
- [28] Stuart Candy and Kelly Kornet. 2019. Turning foresight inside out: An introduction to ethnographic experiential futures. *Journal of Futures Studies* 23, 3 (2019), 3–22. doi:10.6531/JFS.201903_23(3).0002
- [29] Sophie Carruthers, Andrew Pickles, Vicky Slonims, Patricia Howlin, and Tony Charman. 2020. Beyond Intervention into Daily Life: A Systematic Review of Generalisation Following Social Communication Interventions for Young Children with Autism. *Autism Research* 13, 4 (2020), 506–522. doi:10.1002/aur.2264
- [30] EunJeong Cheon and Norman Makoto Su. 2018. Futuristic autobiographies: Weaving participant narratives to elicit values around robots. In *HRI '18*. Association for Computing Machinery, New York, NY, USA, 388–397. doi:10.1145/3171221.3171244
- [31] David Collingridge. 1980. *The Social Control of Technology*. Frances Pinter, London.
- [32] Paul Compensis. 2026. Assessing the integration of social information using AI-generated videos: A brief research report. *Frontiers in Psychology* 16 (2026). doi:10.3389/fpsyg.2025.1505843
- [33] Kathya Cordova-Pozo and Étienne A. J. A. Rouwette. 2023. Types of Scenario Planning and Their Effectiveness: A Review of Reviews. *Futures* 149 (2023), 103153. doi:10.1016/j.futures.2023.103153
- [34] Nicki R. Crick and Kenneth A. Dodge. 1994. A Review and Reformulation of Social Information-Processing Mechanisms in Children's Social Adjustment. *Psychological Bulletin* 115, 1 (1994), 74–101. doi:10.1037/0033-2909.115.1.74
- [35] James J. Cummings and Jeremy N. Bailenson. 2016. How Immersive Is Enough? A Meta-Analysis of the Effect of Immersive Technology on User Presence. *Media Psychology* 19, 2 (2016), 272–309. doi:10.1080/15213269.2015.1015740
- [36] Scott Davidoff, Min Kyung Lee, Anind K. Dey, and John Zimmerman. 2007. Rapidly Exploring Application Design through Speed Dating. In *UbiComp 2007: Ubiquitous Computing*. Springer Berlin Heidelberg, Berlin, Heidelberg, 429–446. doi:10.1007/978-3-540-74853-3_25
- [37] Audrey Desjardins, Cayla Key, Heidi R. Biggs, and Kelsey Aschenbeck. 2019. Bespoke Booklets: A Method for Situated Co-Speculation. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 697–709. doi:10.1145/3322276.3322311
- [38] Mandeep K. Dhami, Lars Wicke, and Dilek Önköl. 2022. Scenario Generation and Scenario Quality Using the Cone of Plausibility. *Futures* 142 (2022), 102995. doi:10.1016/j.futures.2022.102995
- [39] Julia Diemer, Georg W. Alpers, Henrik M. Peperkorn, Youssef Shiban, and Andreas Mühlberger. 2015. The Impact of Perception and Presence on Emotional Reactions: A Review of Research in Virtual Reality. *Frontiers in Psychology* 6 (2015). doi:10.3389/fpsyg.2015.00026
- [40] Christian Dindler and Ole Sejer Iversen. 2007. Fictional Inquiry—Design Collaboration in a Shared Narrative Space. *CoDesign* 3, 4 (2007), 213–234. doi:10.1080/15710880701500187
- [41] Zihan Ding, Chi Jin, Difan Liu, Haitian Zheng, Krishna Kumar Singh, Qiang Zhang, Yan Kang, Zhe Lin, and Yuchen Liu. 2025. DOLLAR: Few-Step Video Generation via Distillation and Latent Reward Optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Piscataway, NJ, USA, 17961–17971. doi:10.1109/ICCV51701.2025.01669
- [42] Benjamin D. Douglas, Patrick J. Ewell, and Markus Brauer. 2023. Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA. *PLOS ONE* 18, 3 (2023), e0279720. doi:10.1371/journal.pone.0279720
- [43] Allison Druin. 2002. The Role of Children in the Design of New Technology. *Behaviour & Information Technology* 21, 1 (2002), 1–25. doi:10.1080/01449290110108659

- [44] Anthony Dunne and Fiona Raby. 2013. *Speculative Everything: Design, Fiction, and Social Dreaming*. MIT Press.
- [45] Chris Elsdén, David Chatting, Abigail C. Durrant, Andrew Garbett, Bettina Nissen, John Vines, and David S. Kirk. 2017. On speculative enactments. In *CHI '17*. Association for Computing Machinery, New York, NY, USA, 5386–5399. doi:10.1145/3025453.3025503
- [46] Hendrik Engelbrecht, Laura Nynke van der Laan, Renske van Enschoot, and Emiel Krahmer. 2022. The Role of Agency and Threat Immediacy in Interactive Digital Narrative Fear Appeals for the Prevention of Excessive Alcohol Use: Randomized Controlled Trial. *JMIR Serious Games* 10, 2 (2022), e32218. doi:10.2196/32218
- [47] Connor Esterwood, Ruijia Guan, Xin Ye, and Lionel P. Robert. 2025. Virtually the Same or Realistically Different? A Meta-Analysis of Real vs. 'Not So Real' Robots. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction*. IEEE, Piscataway, NJ, USA, 559–568. doi:10.1109/HRI61500.2025.10974086
- [48] fal.ai. 2026. MiniMax H3 Max: Faster-than-real-time video generation. Provider benchmark dated 8 September 2026. Accessed 10 September 2026. <https://fal.ai/minimax-h3-max>
- [49] Pedro Gil Farias, Roy Bendor, and Bregje F. van Eekelen. 2022. Social dreaming together: A critical exploration of participatory speculative design. In *PDC '22*. Association for Computing Machinery, New York, NY, USA, 147–154. doi:10.1145/3537797.3537826
- [50] Anna Felnhöfer, Oswald D. Kothgassner, Mareike Schmidt, Anna-Katharina Heinzle, Leon Beutl, Helmut Hlavacs, and Ilse Kryspin-Exner. 2015. Is Virtual Reality Emotionally Arousing? Investigating Five Emotion-Inducing Virtual Park Scenarios. *International Journal of Human-Computer Studies* 82 (2015), 48–56. doi:10.1016/j.ijhcs.2015.05.004
- [51] Batya Friedman and David G. Hendry. 2012. The Envisioning Cards: A Toolkit for Catalyzing Humanistic and Technical Imaginations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1145–1148. doi:10.1145/2207676.2208562
- [52] Tony Fry. 2009. *Design Futuring: Sustainability, Ethics and New Practice*. Berg.
- [53] William Gaver. 2011. Making Spaces: How Design Workbooks Work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1551–1560. doi:10.1145/1978942.1979169
- [54] Friedrich Geiecke and Xavier Jaravel. 2024. Conversations at scale: Robust AI-led interviews with a simple open-source platform. CEPR DP19705; SSRN 4974382. <https://cepr.org/publications/dp19705>. <https://cepr.org/publications/dp19705>
- [55] Parsa Ghaffari and Chris Hokamp. 2025. Narrative Studio: Visual Narrative Exploration Using LLMs and Monte Carlo Tree Search. In *Proceedings of the 7th Workshop on Narrative Understanding*. Association for Computational Linguistics, Albuquerque, New Mexico, 83–96. doi:10.18653/v1/2025.wnu-1.16
- [56] Kristen Gillespie-Lynch, Jennifer B Bisson, Sabine Saade, Rita Obeid, Bella Kofner, Ashley Johnson Harrison, Nidal Daou, Nicholas Tricarico, Jin Delos Santos, William Pinkava, and Allison Jordan. 2022. If You Want to Develop an Effective Autism Training, Ask Autistic Students to Help You. *Autism* 26, 5 (2022), 1082–1094. doi:10.1177/13623613211041006
- [57] Google DeepMind. 2025. Genie 3: A new frontier for world models. <https://deepmind.google/discover/blog/genie-3-a-new-frontier-for-world-models/>. Accessed 6 September 2026. <https://deepmind.google/discover/blog/genie-3-a-new-frontier-for-world-models/>
- [58] Carol A. Gray and Joy D. Garand. 1993. Social Stories: Improving Responses of Students with Autism with Accurate Social Information. *Focus on Autistic Behavior* 8, 1 (1993), 1–10. doi:10.1177/108835769300800101
- [59] Jonathan Green, Tony Charman, Helen McConachie, Catherine Aldred, Vicky Slonims, Patricia Howlin, Ann Le Couteur, Kathy Leadbitter, Kristelle Hudry, Sarah Byford, Barbara Barrett, Kathryn Temple, Wendy Macdonald, and Andrew Pickles. 2010. Parent-Mediated Communication-Focused Treatment in Children with Autism (PACT): A Randomised Controlled Trial. *The Lancet* 375, 9732 (2010), 2152–2160. doi:10.1016/S0140-6736(10)60587-9
- [60] Pam Grossman, Christa Compton, Danielle Igra, Matthew Ronfeldt, Emily Shahan, and Peter W. Williamson. 2009. Teaching Practice: A Cross-Professional Perspective. *Teachers College Record* 111, 9 (2009), 2055–2100. doi:10.1177/016146810911100905
- [61] David Ha and Jürgen Schmidhuber. 2018. World Models. arXiv:1803.10122 <https://arxiv.org/abs/1803.10122>
- [62] Jens Hainmueller, Dominik Hangartner, and Teppei Yamamoto. 2015. Validating vignette and conjoint survey experiments against real-world behavior. *PNAS* 112, 8 (2015), 2395–2400. doi:10.1073/pnas.1416587112
- [63] Hui Han, Siyuan Li, Jiaqi Chen, Yiwen Yuan, Yuling Wu, Yufan Deng, Chak Tou Leong, Hanwen Du, Junchen Fu, Youhua Li, Jie Zhang, Chi Zhang, Li-jia Li, and Yongxin Ni. 2025. Video-Bench: Human-Aligned Video Generation Benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Piscataway, NJ, USA, 18858–18868. doi:10.1109/CVPR52734.2025.01757
- [64] Jongwook Han, Hyeonjin Kim, and Yohan Jo. 2026. RobotValues: Evaluating Household Robots When Human Values Conflict. arXiv:2606.03312 <https://arxiv.org/abs/2606.03312>
- [65] Shaina Herman, Timothy C. Barnum, Paola Emilia Minà, Peter Wozniak, and Jean-Louis van Gelder. 2025. Affect, emotions, and crime decision-making: Emerging insights from immersive 360° video experiments. *Journal of Experimental Criminology* 21, 3 (2025), 965–998. doi:10.1007/s11292-024-09615-y
- [66] Fernanda Herrera, Jeremy Bailenson, Erika Weisz, Elise Ogle, and Jamil Zaki. 2018. Building long-term empathy: A large-scale comparison of traditional and virtual reality perspective-taking. *PLOS ONE* 13, 10 (2018), e0204494. doi:10.1371/journal.pone.0204494
- [67] Arsalan Heydarian, Joao P. Carneiro, David Gerber, Burcin Becerik-Gerber, Timothy Hayes, and Wendy Wood. 2015. Immersive Virtual Environments versus Physical Built Environments: A Benchmarking Study for Building Design and User-Built Environment Explorations. *Automation in Construction* 54 (2015), 116–126. doi:10.1016/j.autcon.2015.03.020

- [68] Botao Amber Hu, Samuel Chua, and Helena Rong. 2026. Protocol Futuring: Speculating Second-Order Dynamics of Protocols in Sociotechnical Infrastructural Futures. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3772318.3790527
- [69] Runsheng Huang, Lara J. Martin, and Chris Callison-Burch. 2024. WHAT-IF: Exploring Branching Narratives by Meta-Prompting Large Language Models. arXiv:2412.10582 <https://arxiv.org/abs/2412.10582>
- [70] Anna E. Jaffe, David DiLillo, Lesa Hoffman, Michelle Haikalas, and Rita E. Dykstra. 2015. Does it hurt to ask? A meta-analysis of participant reactions to trauma research. *Clinical Psychology Review* 40 (2015), 40–56. doi:10.1016/j.cpr.2015.05.004
- [71] Maria Johnson, Aaron M. Tate, Kathleen Tate, Sarah A. Laane, Zhengsi Chang, and Sandra Bond Chapman. 2022. CHARISMA Virtual Social Training: A Digital Health Platform and Protocol. *Frontiers in Virtual Reality* 3 (2022), 1004162. doi:10.3389/frvir.2022.1004162
- [72] Charles M. Judd, Jacob Westfall, and David A. Kenny. 2012. Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology* 103, 1 (2012), 54–69. doi:10.1037/a0028347
- [73] Connie Kasari, Erin Rotheram-Fuller, Jill Locke, and Amanda Gulsrud. 2012. Making the Connection: Randomized Controlled Trial of Social Skills at School for Children with Autism Spectrum Disorders. *Journal of Child Psychology and Psychiatry* 53, 4 (2012), 431–439. doi:10.1111/j.1469-7610.2011.02493.x
- [74] David Kirby. 2010. The Future is Now: Diegetic Prototypes and the Role of Popular Films in Generating Real-World Technological Development. *Social Studies of Science* 40, 1 (2010), 41–70. doi:10.1177/0306312709338325
- [75] Max Knabe, Timothy C. Barnum, Shaina Herman, and Jean-Louis van Gelder. 2026. Focal visual attention in crime decision making: Behavioral insights from immersive 360° video eye-tracking. *Journal of Quantitative Criminology* (2026). doi:10.1007/s10940-026-09657-7
- [76] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuoqiao Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, Kathrina Wu, Qin Lin, Junkun Yuan, Yanxin Long, Aladdin Wang, Andong Wang, Changlin Li, Duoqun Huang, Fang Yang, Hao Tan, Hongmei Wang, Jacob Song, Jiawang Bai, Jianbing Wu, Jinbao Xue, Joey Wang, Kai Wang, Mengyang Liu, Pengyu Li, Shuai Li, Weiyan Wang, Wenqing Yu, Xincheng Deng, Yang Li, Yi Chen, Yutao Cui, Yuanbo Peng, Zhentao Yu, Zhiyu He, Zhiyong Xu, Zixiang Zhou, Zunnan Xu, Yangyu Tao, Qinglin Lu, Songtao Liu, Dax Zhou, Hongfa Wang, Yong Yang, Di Wang, Yuhong Liu, Jie Jiang, and Caesar Zhong. 2024. HunyuanVideo: A Systematic Framework for Large Video Generative Models. arXiv:2412.03603 [cs.CV]
- [77] Diana Kozachek. 2026. Entering the age of hybrid futures: A comparative study of human and GPT-generated scenarios. *Futures* 176 (2026), 103705. doi:10.1016/j.futures.2025.103705
- [78] Sandjar Kozubae, Chris Elsdon, Noura Howell, Marie Louise Juul Søndergaard, Nick Merrill, Britta Schulte, and Richmond Y. Wong. 2020. Expanding Modes of Reflection in Design Futuring. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3313831.3376526
- [79] Saskia F. Kuliga, Tyler Thrash, Ruth C. Dalton, and Christoph Hölscher. 2015. Virtual Reality as an Empirical Research Tool: Exploring User Experience in a Real Building and a Corresponding Virtual Model. *Computers, Environment and Urban Systems* 54 (2015), 363–375. doi:10.1016/j.compenvurbsys.2015.09.006
- [80] Nilay Kumar, Priyansh Bhandari, and G. Maragatham. 2025. Bridging Text and Video Generation: A Survey. *arXiv preprint* (2025). doi:10.48550/arXiv.2510.04999
- [81] Gary W. Ladd and Jacquelyn Mize. 1983. A Cognitive-Social Learning Model of Social-Skill Training. *Psychological Review* 90, 2 (1983), 127–157. doi:10.1037/0033-295X.90.2.127
- [82] Kathy Leadbitter, Karen Leneh Buckle, Ceri Ellis, and Martijn Dekker. 2021. Autistic Self-Advocacy and the Neurodiversity Movement: Implications for Autism Early Intervention Research and Practice. *Frontiers in Psychology* 12 (2021), 635690. doi:10.3389/fpsyg.2021.635690
- [83] Jorge Leandro, Sudha Rao, Michael Xu, Weijia Xu, Nebojsa Jojic, Chris Brockett, and Bill Dolan. 2024. GENEVA: Generating and visualizing branching narratives using LLMs. In *2024 IEEE Conference on Games (CoG)*. IEEE, Piscataway, NJ, USA, 1–5. arXiv:2311.09213 doi:10.1109/COG60054.2024.10645625
- [84] Jiwon Lee, Yewon Kim, Chang-Min Kim, Suhyun Lim, Chang Hyun Park, Subin Cho, and Tek-Jin Nam. 2026. Reminiscences of Futures (RoF): Designing and studying encounters with embodied, AI-generated future selves. In *CHI '26*. Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/3772318.3791750
- [85] Jamy Li. 2015. The Benefit of Being Physically Present: A Survey of Experimental Works Comparing Copresent Robots, Telepresent Robots and Virtual Agents. *International Journal of Human-Computer Studies* 77 (2015), 23–37. doi:10.1016/j.ijhcs.2015.01.001
- [86] Joseph Lindley and Paul Coulton. 2015. Back to the future: 10 years of design fiction. In *British HCI '15*. Association for Computing Machinery, New York, NY, USA, 210–211. doi:10.1145/2783446.2783592
- [87] Huan Ling, Seung Wook Kim, Antonio Torralba, Sanja Fidler, and Karsten Kreis. 2024. Align Your Gaussians: Text-to-4D with Dynamic 3D Gaussians and Composed Diffusion Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Piscataway, NJ, USA, 8576–8588. arXiv:2312.13763 doi:10.1109/CVPR52733.2024.00819
- [88] Jinxiu Liu, Xuanming Liu, Kangfu Mei, Yandong Wen, Ming-Hsuan Yang, and Weiyang Liu. 2026. Streaming Autoregressive Video Generation via Diagonal Distillation. *arXiv preprint* (2026). doi:10.48550/arXiv.2603.09488
- [89] Joan Llobera, Alejandro Beacco, Ramon Oliva, Gizem Şenel, Domna Banakou, and Mel Slater. 2021. Evaluating participant responses to a virtual reality experience using reinforcement learning. *Royal Society Open Science* 8, 9 (2021), 210537. doi:10.1098/rsos.210537

- [90] Rachel Lowy, Lan Gao, Kaely Hall, and Jennifer G Kim. 2023. Toward Inclusive Mindsets: Design Opportunities to Represent Neurodivergent Work Experiences to Neurotypical Co-Workers in Virtual Reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3544548.3581399
- [91] Sasha Luccioni, Christopher Akiki, Margaret Mitchell, and Yacine Jernite. 2023. Stable Bias: Evaluating societal representations in diffusion models. In *NeurIPS 2023 Datasets & Benchmarks*. https://proceedings.neurips.cc/paper_files/paper/2023/hash/b01153e7112b347d8ed54f317840d8af-Abstract-Datasets_and_Benchmarks.html
- [92] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. 2024. Dynamic 3D Gaussians: Tracking by Persistent Dynamic View Synthesis. In *International Conference on 3D Vision (3DV)*. IEEE, Piscataway, NJ, USA, 800–809. arXiv:2308.09713 doi:10.1109/3DV62453.2024.00044
- [93] Laura Lundy. 2007. ‘Voice’ Is Not Enough: Conceptualising Article 12 of the United Nations Convention on the Rights of the Child. *British Educational Research Journal* 33, 6 (2007), 927–942. doi:10.1080/01411920701657033
- [94] Wendy E. Mackay and Anne Laure Fayard. 1999. Video Brainstorming and Prototyping: Techniques for Participatory Design. In *CHI ’99 Extended Abstracts on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 118–119. doi:10.1145/632716.632790
- [95] Ifigeneia Mavridou, Emili Balaguer-Ballester, Charles Nduka, and Ellen Seiss. 2023. A Reliable and Robust Online Validation Method for Creating a Novel 3D Affective Virtual Environment and Event Library (AVEL). *PLOS ONE* 18, 4 (2023), e0278065. doi:10.1371/journal.pone.0278065
- [96] Sveta Mayer, Iroise Dumontheil, Hannah R. Wilkinson, Kaska Porayska-Pomsta, Emily K. Farran, Andrew K. Tolmie, Denis Mareschal, and the UnLocke Team. 2025. SEE+ Computerized Classroom-Based Training Enhances 7- to 10-Year-Olds’ Socio-Emotional Cognition through Observation and Inference. *PLOS ONE* 20, 9 (2025), e0330934. doi:10.1371/journal.pone.0330934
- [97] Joshua McVeigh-Schultz, Max Kreminski, Keshav Prasad, Perry Hoberman, and Scott S. Fisher. 2018. Immersive design fiction: Using VR to prototype speculative interfaces and interaction rituals within a virtual storyworld. In *DIS ’18*. Association for Computing Machinery, New York, NY, USA, 817–829. doi:10.1145/3196709.3196793
- [98] Amy Meenaghan, Claire Nee, Jean-Louis van Gelder, Marco Otte, and Zarah Vernham. 2018. Getting closer to the action: Using the virtual enactment method to understand burglary. *Deviant Behavior* 39, 4 (2018), 437–460. doi:10.1080/01639625.2017.1407104
- [99] Jamie N. Mikeska, Heather Howell, and Devon Kinsey. 2022. Examining the usability and viability of using a simulated classroom environment to prepare preservice science teachers during and after the COVID-19 pandemic. *Disciplinary and Interdisciplinary Science Education Research* 4, 1, Article 23 (2022). doi:10.1186/s43031-022-00054-1
- [100] Damian E. M. Milton. 2012. On the Ontological Status of Autism: The ‘Double Empathy Problem’. *Disability & Society* 27, 6 (2012), 883–887. doi:10.1080/09687599.2012.710008
- [101] MiniMax. 2026. MiniMax H3: An open model breaking the boundaries between tasks and modalities. Released 31 July 2026. Accessed 6 September 2026. <https://www.minimax.io/blog/minimax-h3>
- [102] Michael J. Muller. 2007. Participatory Design: The Third Space in HCI. In *The Human-Computer Interaction Handbook* (2 ed.). CRC Press, 1061–1081. doi:10.1201/9781410615862.ch54
- [103] Lala Muradova and Edana Beauvais. 2025. Promoting pro-environmental beliefs and behaviour: Choose-your-own story futuristic climate game. *PLOS ONE* 20, 3 (2025), e0317773. doi:10.1371/journal.pone.0317773
- [104] Bilge Mutlu, Toshiyuki Shiwa, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2009. Footing in Human-Robot Conversations: How Robots Might Shape Participant Roles Using Gaze Cues. In *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction*. Association for Computing Machinery, New York, NY, USA, 61–68. doi:10.1145/1514095.1514109
- [105] Mohammad Nadeem, Shahab Saquib Sohail, Erik Cambria, Björn W. Schuller, and Amir Hussain. 2025. Gender bias in text-to-video generation models: A case study of Sora. *IEEE Intelligent Systems* 40, 3 (2025), 10–15. arXiv:2501.01987 doi:10.1109/MIS.2025.3561475
- [106] Michelle R. Nario-Redmond, Dobromir Gospodinov, and Angela Cobb. 2017. Crip for a Day: The Unintended Negative Consequences of Disability Simulations. *Rehabilitation Psychology* 62, 3 (2017), 324–333. doi:10.1037/rep0000127
- [107] Lisa P. Nathan, Predrag V. Klasnja, and Batya Friedman. 2007. Value Scenarios: A Technique for Envisioning Systemic Effects of New Technologies. In *CHI ’07 Extended Abstracts on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 2585–2590. doi:10.1145/1240866.1241046
- [108] Michael Nebeling and Maximilian Speicher. 2018. The Trouble with Augmented Reality/Virtual Reality Authoring Tools. In *2018 IEEE International Symposium on Mixed and Augmented Reality Adjunct*. IEEE, Piscataway, NJ, USA, 333–337. doi:10.1109/ISMAR-Adjunct.2018.00098
- [109] Renee Noortman, Britta F. Schulte, Paul Marshall, Saskia Bakker, and Anna L. Cox. 2019. HawkEye: Deploying a Design Fiction Probe. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 422, 14 pages. doi:10.1145/3290605.3300652
- [110] William Odom, John Zimmerman, Scott Davidoff, Jodi Forlizzi, Anind K. Dey, and Min Kyung Lee. 2012. A fieldwork of the future with user enactments. In *DIS ’12*. Association for Computing Machinery, New York, NY, USA, 338–347. doi:10.1145/2317956.2318008
- [111] Hyangkyeong Oh, Narae Hong, Kwanguk Kim, Dong-Hyun Oh, Jae-Jin Kim, and Eunjo Kim. 2026. Stepping into Virtual School: A Pilot Study of VR-Based Social Skills Training Integrating Emotion- and Knowledge-Based Approaches. *Virtual Reality* 30, 1 (2026), 40. doi:10.1007/s10055-025-01285-z
- [112] Wyatt Olson and James Pierce. 2025. Exploring the Use of Speculative Concept Films for Co-Speculation around Data Ethics. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 235, 21 pages. doi:10.1145/3706598.3713245

- [113] Antti Oulasvirta, Esko Kurvinen, and Tomi Kankainen. 2003. Understanding contexts by being there: case studies in bodystorming. *Personal and Ubiquitous Computing* 7, 2 (2003), 125–134. doi:10.1007/s00779-003-0238-7
- [114] Ana Paiva, João Dias, Daniel Sobral, Ruth Aylett, Sarah Woods, Lynne Hall, and Carsten Zoll. 2005. Learning by Feeling: Evoking Empathy with Synthetic Characters. *Applied Artificial Intelligence* 19, 3–4 (2005), 235–266. doi:10.1080/08839510590910165
- [115] Anastasia Panfilova, Vadim Bolshev, Mikhail Mozikov, Vladislav Latynov, Alexander Vanin, Timofei Nestik, Veronika Nourkova, Arina Vlasova, Mikhail Kozin, Aleksandr Serohvostov, Elizaveta Tarasova, and Sergey Nikolenko. 2026. The AI interviewer: Multi-faceted evaluation of adaptive questioning by large language models. *Scientific Reports* 16, 1 (2026). doi:10.1038/s41598-026-46517-7
- [116] Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 1–22. doi:10.1145/3586183.3606763
- [117] Pat Pataranutaporn, Kavin Winson, Peggy Yin, Auttasak Lapapirojn, Pichayoot Ouppaphan, Monchai Lertsutthiwong, Pattie Maes, and Hal E. Herschfield. 2024. Future You: A conversation with an AI-generated future self reduces anxiety, negative emotions, and increases future self-continuity. In *2024 IEEE Frontiers in Education Conference (FIE)*. IEEE, Piscataway, NJ, USA, 1–10. arXiv:2405.12514 doi:10.1109/FIE61694.2024.10893443
- [118] William Peebles and Saining Xie. 2023. Scalable Diffusion Models with Transformers. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Piscataway, NJ, USA, 4172–4182. doi:10.1109/ICCV51070.2023.00387
- [119] Jin Peng, Jiyue Li, Danyi Li, Yihua Fang, Chi Zhang, Mark W. Fraser, and Shenyang Guo. 2024. Systematic Review and Meta-Analysis of Interventions to Improve Children’s Social Information-Processing Skills. *Research on Social Work Practice* 34, 5 (2024), 507–534. doi:10.1177/10497315231182449
- [120] Léa Pessin and Kevin Munger. 2026. Beyond text: Using AI-generated visual conjoints to study gender and housework attribution. *Sociological Science* 13 (2026), 661–684. doi:10.15195/v13.a26
- [121] Melina Petsolari, Seray Ibrahim, and Petr Slovak. 2024. Socio-technical Imaginaries: Envisioning and Understanding AI Parenting Supports through Design Fiction. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 98, 27 pages. doi:10.1145/3613904.3642619
- [122] Andrew Pickles, Ann Le Couteur, Kathy Leadbitter, Erica Salomone, Rachel Cole-Fletcher, Hannah Tobin, Isobel Gammer, Jessica Lowry, George Vamvakas, Sarah Byford, Catherine Aldred, Vicky Slonims, Helen McConachie, Patricia Howlin, Jeremy R. Parr, Tony Charman, and Jonathan Green. 2016. Parent-Mediated Social Communication Therapy for Young Children with Autism (PACT): Long-Term Follow-Up of a Randomised Controlled Trial. *The Lancet* 388, 10059 (2016), 2501–2509. doi:10.1016/S0140-6736(16)31229-6
- [123] James Pierce. 2021. In Tension with Progression: Grasping the Frictional Tendencies of Speculative, Critical, and Other Alternative Designs. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 617, 19 pages. doi:10.1145/3411764.3445406
- [124] Rachel Poonsriwong, Chayapat Archiwangruek, Constanze Albrecht, Peggy Yin, Nattavudh Powdthavee, Hal Herschfield, Monchai Lertsutthiwong, Kavin Winson, and Pat Pataranutaporn. 2026. Simulating life paths with digital twins: AI-generated future selves influence decision-making and expand human choice. In *Proceedings of the Augmented Humans International Conference 2026*. Association for Computing Machinery, New York, NY, USA, 264–278. arXiv:2512.05397 doi:10.1145/3795011.3797152
- [125] David Porfirio, Evan Fisher, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2019. Bodystorming human-robot interactions. In *UIST ’19*. Association for Computing Machinery, New York, NY, USA, 479–491. doi:10.1145/3332165.3347957
- [126] Isabella Pu, Kantwon Rogers, Linh Dieu Dinh, Sharifa Alghowinem, and Cynthia Breazeal. 2026. Who’s the Boss? Children Negotiate Robot Control across Role and Context. In *Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction*. Association for Computing Machinery, New York, NY, USA, 395–405. doi:10.1145/3757279.3785548
- [127] Iyad Rahwan, Azim Shariff, and Jean-François Bonnefon. 2025. The science fiction science method. *Nature* 644, 8075 (2025), 51–58. arXiv:2508.03430. doi:10.1038/s41586-025-09194-6
- [128] Rafael Ramirez and Cynthia Selin. 2014. Plausibility and Probability in Scenario Planning. *Foresight* 16, 1 (2014), 54–74. doi:10.1108/FS-08-2012-0061
- [129] Shannon L. Rawski, Joshua R. Foster, and Jeremy Bailenson. 2022. Sexual Harassment Bystander Training Effectiveness: Experimentally Comparing 2D Video to Virtual Reality Practice. *Technology, Mind, and Behavior* 3, 2 (2022), 302–321. doi:10.1037/tmb0000074
- [130] Linda M. Reeves, Rebecca M. Giles, and Todd Johnson. 2019. Exploring Preservice Teachers’ Perceptions of Preparedness to Teach Students with Disabilities. *Journal of the American Academy of Special Education Professionals* (2019), 112–129. doi:10.64546/jaasep.398
- [131] Jiawei Ren, Liang Pan, Jiaxiang Tang, Chi Zhang, Ang Cao, Gang Zeng, and Ziwei Liu. 2023. DreamGaussian4D: Generative 4D Gaussian Splatting. arXiv:2312.17142 <https://arxiv.org/abs/2312.17142>
- [132] Xuanchi Ren, Tianchang Shen, Jiahui Huang, Huan Ling, Yifan Lu, Merlin Nimier-David, Thomas Müller, Alexander Keller, Sanja Fidler, and Jun Gao. 2025. GEN3C: 3D-Informed World-Consistent Video Generation with Precise Camera Control. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Piscataway, NJ, USA, 6121–6132. arXiv:2503.03751 doi:10.1109/CVPR52734.2025.00574
- [133] Laurel Riek. 2012. Wizard of Oz studies in HRI: A systematic review and new reporting guidelines. *Journal of Human-Robot Interaction* 1, 1 (2012), 119–136. doi:10.5898/JHRI.1.1.Riek
- [134] Ronda Ringfort-Felner, Judith Dörrenbächer, and Marc Hassenzahl. 2025. The Quality of Speculation—A Scoping Review. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 2373–2394. doi:10.1145/3715336.3735794
- [135] Rebecca Roberts, James Stacey, Susannah Jenner, and Eloise Maguire. 2023. Are Extended Reality Interventions Effective in Helping Autistic Children to Enhance Their Social Skills? A Systematic Review. *Review Journal of Autism and Developmental Disorders* 10, 4 (2023), 729–748.

- doi:10.1007/s40489-022-00320-y
- [136] Chandler G. Robinson, Ian T. Adams, Matthew W. Logan, and J. Pete Blair. 2026. AI-generated human stimuli for experimental social science. *Journal of Experimental Criminology* (2026). doi:10.1007/s11292-026-09731-x
- [137] Eileen Roesler, Dietrich Manzey, and Linda Onnasch. 2023. Embodiment matters in social HRI research: Effectiveness of anthropomorphism on subjective and objective outcomes. *ACM Transactions on Human-Robot Interaction* 12, 1 (2023), 1–9. doi:10.1145/3555812
- [138] Peter H. Rossi and Steven L. Nock. 1982. *Measuring social judgments: The factorial survey approach*. Sage. https://books.google.com/books/about/Measuring_Social_Judgments.html?id=pai6AAAAIAAJ
- [139] Christian Roth and Hartmut Koenitz. 2016. Evaluating the user experience of interactive digital narrative. In *Proceedings of the 1st International Workshop on Multimedia Alternate Realities*. Association for Computing Machinery, New York, NY, USA, 31–36. doi:10.1145/2983298.2983302
- [140] Aitor Rovira and Mel Slater. 2022. Encouraging bystander helping behaviour in a violent incident: A virtual reality study using reinforcement learning. *Scientific Reports* 12, 1 (2022), 3843. doi:10.1038/s41598-022-07872-3
- [141] Aitor Rovira, Richard Southern, David Swapp, Claire Campbell, Jian J. Zhang, Mark Levine, and Mel Slater. 2021. Bystander affiliation influences intervention behavior: A virtual reality study. *SAGE Open* 11, 3 (2021). doi:10.1177/21582440211040076
- [142] Antti Salovaara and Leevi Vahvelainen. 2025. Triangulating on possible futures: Conducting user studies on several futures instead of only one. In *CHI '25*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3706598.3713565
- [143] Elizabeth B.-N. Sanders and Pieter Jan Stappers. 2008. Co-creation and the new landscapes of design. *CoDesign* 4, 1 (2008), 5–18. doi:10.1080/15710880701875068
- [144] Melanie A. Sarge, Hark-Shin Kim, and John A. Velez. 2020. An Anti-Sim Intervention: The Role of Perspective Taking in Combating Public Stigma with Virtual Simulations. *Cyberpsychology, Behavior, and Social Networking* 23, 1 (2020), 41–51. doi:10.1089/cyber.2019.0678
- [145] Dennis Schleicher, Peter Jones, and Oksana Kachur. 2010. Bodystorming as Embodied Designing. *Interactions* 17, 6 (2010), 47–51. doi:10.1145/1865245.1865256
- [146] Douglas Schuler and Aki Namioka (Eds.). 2017. *Participatory Design: Principles and Practices*. CRC Press. Originally published 1993. doi:10.1201/9780203744338
- [147] Nicole Schwitter. 2025. Using artificial intelligence to generate visual vignettes in factorial survey experiments. *Social Science Computer Review* (2025). doi:10.1177/08944393251392916
- [148] Adalberto L. Simeone, Robbe Cools, Stan Depuydt, João Maria Gomes, Piet Goris, Joseph Grocott, Augusto Esteves, and Kathrin Gerling. 2022. Immersive speculative enactments: Bringing future scenarios and technology to life using virtual reality. In *CHI '22, Article 17*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3491102.3517492
- [149] Jesper Simonsen and Toni Robertson (Eds.). 2013. *Routledge International Handbook of Participatory Design*. Routledge, New York, NY. doi:10.4324/9780203108543
- [150] Mel Slater. 2009. Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Philosophical Transactions of the Royal Society B* 364, 1535 (2009), 3549–3557. doi:10.1098/rstb.2009.0138
- [151] Mel Slater, Angus Antley, Adam Davison, David Swapp, Christoph Guger, Chris Barker, Nancy Pistrang, and Maria V. Sanchez-Vives. 2006. A virtual reprise of the Stanley Milgram obedience experiments. *PLOS ONE* 1, 1 (2006), e39. doi:10.1371/journal.pone.0000039
- [152] Mel Slater, Aitor Rovira, Richard Southern, David Swapp, Jian J. Zhang, Claire Campbell, and Mark Levine. 2013. Bystander responses to a violent incident in an immersive virtual environment. *PLOS ONE* 8, 1 (2013), e52766. doi:10.1371/journal.pone.0052766
- [153] Susan H. Spence. 2003. Social Skills Training with Children and Young People: Theory, Evidence and Practice. *Child and Adolescent Mental Health* 8, 2 (2003), 84–96. doi:10.1111/1475-3588.00051
- [154] Hongyi Tao and Dhaval Vyas. 2025. "Housing diversity means diverse housing": Blending generative AI into speculative design in rural co-housing communities. In *CHI '25*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3706598.3713906
- [155] David Thue, Vadim Bulitko, Marcia Spetch, and Trevon Romanuik. 2010. Player agency and the relevance of decisions. In *ICIDS 2010, LNCS 6432*. Springer Berlin Heidelberg, Berlin, Heidelberg, 210–215. doi:10.1007/978-3-642-16638-9_26
- [156] Bruce Tognazzini. 1994. The "Starfire" Video Prototype Project: A Case History. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 99–105. doi:10.1145/191666.191712
- [157] Nathan Tsoi, Mohamed Hussein, Olivia Fugikawa, J. D. Zhao, and Marynel Vázquez. 2021. An approach to deploy interactive robotic simulators on the web for HRI experiments: Results in social robot navigation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Piscataway, NJ, USA, 7528–7535. doi:10.1109/IROS51168.2021.9636319
- [158] Nathan Tsoi, Rachel Sterneck, Xuan Zhao, and Marynel Vázquez. 2024. Influence of simulation and interactivity on human perceptions of a robot during navigation tasks. *ACM Transactions on Human-Robot Interaction* 13, 4 (2024), 1–19. doi:10.1145/3675784
- [159] Jakob Uhl, Marita Huber, Aysenur Gurel, Sophie Westfahl, Stefan Killian, Hannes Schweiger, Alexander Rind, Georg Regal, and Manfred Tscheligi. 2026. Branching Foresight: A novel interaction concept for AI-generated scenario exploration. In *CHI '26, Article 1272*. Association for Computing Machinery, New York, NY, USA, 1–24. doi:10.1145/3772318.3793527
- [160] Jean-Louis van Gelder, Timothy C. Barnum, Shaina Herman, and Peter Wozniak. 2026. The MAXLab aggression and bystander intervention scenario set (MAXLab_ABISS): A modular scenario set for studying decision making in situations of interpersonal violence in virtual reality. *Journal of Experimental Criminology* 22, 1 (2026), 249–269. doi:10.1007/s11292-024-09645-6

- [161] Jean-Louis van Gelder, Reinout E. de Vries, Andrew Demetriou, Iris van Sintemaartensdijk, and Tara Donker. 2019. The Virtual Reality Scenario Method: Moving from imagination to immersion in criminal decision-making research. *Journal of Research in Crime and Delinquency* 56, 3 (2019), 451–480. doi:10.1177/0022427818819696
- [162] Iris van Sintemaartensdijk, Jean-Louis van Gelder, Jan-Willem van Prooijen, Claire Nee, Marco Otte, and Paul van Lange. 2021. Mere Presence of Informal Guardians Deters Burglars: A Virtual Reality Study. *Journal of Experimental Criminology* 17, 4 (2021), 657–676. doi:10.1007/s11292-020-09430-1
- [163] Anand P. A. van Zelderen, Theodore C. Masters-Waage, Nicky Dries, Jochen I. Menges, and Diana R. Sanchez. 2025. Simulating virtual organizations for research: A comparative empirical evaluation of text-based, video, and virtual reality video vignettes. *Organizational Research Methods* 28, 3 (2025), 457–486. doi:10.1177/10944281241246770
- [164] Yamil Ricardo Velez. 2025. Crowdsourced adaptive surveys. *Political Analysis* 33, 4 (2025), 284–297. doi:10.1017/pan.2024.34
- [165] Yamil Ricardo Velez and Patrick Liu. 2025. Confronting core issues: A critical assessment of attitude polarization using tailored experiments. *American Political Science Review* 119, 2 (2025), 1036–1053. doi:10.1017/S0003055424000819
- [166] Rogier E. J. Verhoef, Anouk van Dijk, Esmée E. Verhulp, and Bram O. de Castro. 2021. Interactive Virtual Reality Assessment of Aggressive Social Information Processing in Boys with Behaviour Problems: A Pilot Study. *Clinical Psychology & Psychotherapy* 28, 3 (2021), 489–499. doi:10.1002/cpp.2620
- [167] Laurie Vertelney. 1989. Using Video to Prototype User Interfaces. *ACM SIGCHI Bulletin* 21, 2 (1989), 57–61. doi:10.1145/70609.70615
- [168] Joseph Voros. 2003. A generic foresight process framework. *Foresight* 5, 3 (2003), 10–21. doi:10.1108/14636680310698379
- [169] Ron Wakkary, Doenja Oogjes, and Armi Behzad. 2022. Two years or more of co-speculation: Polylogues of philosophers, designers, and a tilting bowl. *ACM Transactions on Computer-Human Interaction* 29, 5 (2022), 1–44. doi:10.1145/3514235
- [170] Sara Walton, Paula O’Kane, and Diane Ruwhiu. 2019. Developing a Theory of Plausibility in Scenario Building: Designing Plausible Scenarios. *Futures* 111 (2019), 42–56. doi:10.1016/j.futures.2019.03.002
- [171] Wan Team, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wenten Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. 2025. Wan: Open and Advanced Large-Scale Video Generative Models. arXiv:2503.20314 [cs.CV]
- [172] Angelina Wang, Jamie Morgenstern, and John P. Dickerson. 2025. Large Language Models That Replace Human Participants Can Harmfully Misportray and Flatten Identity Groups. *Nature Machine Intelligence* 7, 3 (2025), 400–411. doi:10.1038/s42256-025-00986-z
- [173] Qianqian Wang, Vickie Ye, Hang Gao, Weijia Zeng, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. 2025. Shape of Motion: 4D Reconstruction from a Single Video. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Piscataway, NJ, USA, 9660–9672. arXiv:2407.13764 doi:10.1109/ICCV51701.2025.00901
- [174] Yi Wang, John Joon Young Chung, Melissa Roemmele, Yuqian Sun, Tiffany Wang, Shm Garanganao Almeda, Brett A. Halperin, Yuwen Lu, and Max Kreminski. 2026. Elsewise: Authoring AI-Based Interactive Narrative with Possibility Space Visualization. arXiv:2601.15295 <https://arxiv.org/abs/2601.15295>
- [175] Jacob Westfall, Charles M. Judd, and David A. Kenny. 2015. Replicating studies in which samples of participants respond to samples of stimuli. *Perspectives on Psychological Science* 10, 3 (2015), 390–399. doi:10.1177/1745691614564879
- [176] Richmond Y. Wong, Deirdre K. Mulligan, Ellen Van Wyk, James Pierce, and John Chuang. 2017. Eliciting Values Reflections by Engaging Privacy Futures Using Design Workbooks. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW, Article 111 (2017), 27 pages. doi:10.1145/3134746
- [177] Richmond Y. Wong, Jason Caleb Valdez, Ashten Alexander, Ariel Chiang, Olivia Quesada, and James Pierce. 2023. Broadening Privacy and Surveillance: Eliciting Interconnected Values with a Scenarios Workbook on Smart Home Cameras. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 1093–1113. doi:10.1145/3563657.3596012
- [178] World Labs. 2026. Atlas: A World Model for Spatial Intelligence. <https://www.worldlabs.ai/blog/atlas>. Accessed 6 September 2026. <https://www.worldlabs.ai/blog/atlas>
- [179] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 2024. 4D Gaussian Splatting for Real-Time Dynamic Scene Rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Piscataway, NJ, USA, 20310–20320. arXiv:2310.08528 doi:10.1109/CVPR52733.2024.01920
- [180] Rundi Wu, Ruiqi Gao, Ben Poole, Alex Trevithick, Changxi Zheng, Jonathan T. Barron, and Aleksander Holynski. 2025. CAT4D: Create Anything in 4D with Multi-View Video Diffusion Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Piscataway, NJ, USA, 26057–26068. arXiv:2411.18613 doi:10.1109/CVPR52734.2025.02427
- [181] Yuning Wu, Bingheng Ryan Qiu, and Spencer Swadish. 2026. Conducting online randomized experiments: Lessons from a policing survey using AI-edited videos. *Journal of Experimental Criminology* (2026). doi:10.1007/s11292-025-09722-4
- [182] Alexander Wuttke, Matthias Aßenmacher, Christopher Klammer, Max M. Lang, Quirin Würschinger, and Frauke Kreuter. 2025. AI conversational interviewing: Transforming surveys with LLMs as adaptive interviewers. In *Proceedings of the 9th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature (LaTeCH-CLfL 2025)*. Association for Computational Linguistics, Manuscript submitted to ACM

- Albuquerque, New Mexico, 179–204. arXiv:2410.01824 doi:10.18653/v1/2025.latechclfl-1.17
- [183] Alexander Wuttke, Max Melchior Lang, Christopher Klammer, Quirin Würschinger, and Frauke Kreuter. 2026. AI conversational interviewing: Scaling up semi-structured and in-depth interviews. arXiv:2606.20064 <https://arxiv.org/abs/2606.20064>
- [184] Michael F. Xu, Enhui Zhao, Yawen Zhang, Joseph E. Michaelis, Sarah Sebo, and Bilge Mutlu. 2026. Designing Robots for Families: In-Situ Prototyping for Contextual Reminders on Family Routines. In *Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction*. Association for Computing Machinery, New York, NY, USA, 356–365. doi:10.1145/3757279.3788654
- [185] Shuchang Xu, Xiaofu Jin, Wenshuo Zhang, Huamin Qu, and Yukang Yan. 2025. Branch Explorer: Leveraging branching narratives to support interactive 360° video viewing for blind and low vision users. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 1–18. arXiv:2507.09959 doi:10.1145/3746059.3747791
- [186] Zhen Xu, Zhengqin Li, Zhao Dong, Xiaowei Zhou, Richard Newcombe, and Zhaoyang Lv. 2025. 4DGT: Learning a 4D Gaussian Transformer Using Real-World Monocular Videos. In *Advances in Neural Information Processing Systems (NeurIPS)*. Neural Information Processing Systems Foundation, Inc., 47576–47602. Spotlight. arXiv:2506.08015 doi:10.52202/085713-1419
- [187] Xiao Xue, Xinyang Li, Boyang Jia, Jiachen Du, and Xinyi Fu. 2024. Who Should Hold Control? Rethinking Empowerment in Home Automation among Cohabitants through the Lens of Co-Design. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3613904.3642866
- [188] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, Da Yin, Yuxuan Zhang, Weihang Wang, Yean Cheng, Bin Xu, Xiaotao Gu, Yuxiao Dong, and Jie Tang. 2024. CogVideoX: Text-to-Video Diffusion Models with an Expert Transformer. *arXiv preprint* (2024). doi:10.48550/arXiv.2408.06072
- [189] Tianwei Yin, Qiang Zhang, Richard Zhang, William T. Freeman, Fredo Durand, Eli Shechtman, and Xun Huang. 2025. From Slow Bidirectional to Fast Autoregressive Video Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Piscataway, NJ, USA, 22963–22974. doi:10.1109/CVPR52734.2025.02138
- [190] Xinyan Yu, Marius Hoggemüller, Tram Thi Minh Tran, and Martin Tomitsch. 2026. Fostering Design–Policy Collaboration through Contestation: An Adversarial Futuring Method. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3772363.3799303
- [191] Hamza Zarglayoun, Juliette Laurendeau-Martin, Ange Tato, Evelyn Vera-Estay, Aurélie Blondin, Arnaud Lamy-Brunelle, Sameh Chaieb, Frédéric Morasse, Aude Dufresne, Roger Nkambou, and Miriam H. Beauchamp. 2022. Assessing and Optimizing Socio-Moral Reasoning Skills: Findings from the MorALERT Serious Video Game. *Frontiers in Psychology* 12 (2022), 767596. doi:10.3389/fpsyg.2021.767596
- [192] Jichen Zhu, Connor M. Kerns, James Connell, and Natalie Lyon. 2016. Using Interactive Social Story Games to Teach Social Skills to Children with Autism. In *Proceedings of the First Joint International Conference of DiGRA and FDG*. Digital Games Research Association. doi:10.26503/dl.v2016i1.783
- [193] John Zimmerman. 2005. Video Sketches: Exploring Pervasive Computing Interaction Designs. *IEEE Pervasive Computing* 4, 4 (2005), 91–94. doi:10.1109/MPRV.2005.91