

Storycaster: An AI System for Immersive Room-based Storytelling

Naisha Agarwal*
University of California, Los Angeles
Los Angeles, California, USA
naishaa@g.ucla.edu

Judith Amores
Microsoft Research
Cambridge, Massachusetts, USA
judithamores@microsoft.com

Andrew D Wilson
Microsoft Research
Redmond, Washington, USA
awilson@microsoft.com



Figure 1: *Storycaster* transforms your room into a responsive storytelling environment. Multiple cameras and projectors are used to augment the walls and furniture of the room, setting the scene for a story that is told through dialogue, ambient sound effects, and music. Here we see three examples of automatically generated scenes: a wood-paneled room with overstuffed leather chairs, a fantasy-inspired view into space, and a witch's cottage. Depth sensing and projection mapping are used to adapt visual content to the room's geometry and surfaces.

Abstract

While Cave Automatic Virtual Environment (CAVE) systems have long enabled room-scale virtual reality and various kinds of interactivity, their content has largely remained predetermined. We present *Storycaster*, a generative AI CAVE system that transforms physical rooms into responsive storytelling environments. Unlike headset-based VR, *Storycaster* preserves spatial awareness, using live camera feeds to augment the walls with cylindrical projections, allowing users to create worlds that blend with their physical surroundings. Additionally, our system enables object-level editing, where physical items in the room can be transformed to their virtual counterparts in a story. A narrator agent guides participants, enabling them to co-create stories that evolve in response to voice commands, with each scene enhanced by generated ambient audio, dialogue, and imagery. Participants in our study ($n = 13$) found the system highly immersive and engaging, identifying the narrator and audio as the most impactful elements, while also highlighting areas of improvement in latency and image resolution.

CCS Concepts

• **Human-centered computing** → **Mixed / augmented reality**; *Natural language interfaces*; User studies.

*This work was done during an internship at Microsoft Research.



This work is licensed under a Creative Commons Attribution 4.0 International License.
CHI '26, Barcelona, Spain
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2278-3/26/04
<https://doi.org/10.1145/3772318.3791305>

Keywords

Immersive Storytelling, Projection Mapping, Narratives, Human-AI co-creation, agentic AI

ACM Reference Format:

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 (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 26 pages. <https://doi.org/10.1145/3772318.3791305>

1 Introduction

Spatial Augmented Reality (SAR) systems use precise geometric models of physical environments to seamlessly blend rendered virtual spaces into the real world. Systems such as Room Alive [30] demonstrate how projection mapping techniques can transform a room without needing head-mounted displays (HMDs), allowing users to remain present in and directly interact with their surroundings. However, these kinds of experiences have traditionally relied on pre-authored content that cannot adapt dynamically to the user's physical environment or creative goals. Authoring such experiences is difficult since designers rarely have access to the user's room, yet convincing projection mapping requires content that is highly situated and spatially aligned with its environment.

Recently developed AI image generation techniques offer a new opportunity to address this challenge. With the advent of large language models, we can consider the automatic generation of not just the visuals but rather the entire interaction itself, from different environments, characters, and narrative structure, all responding directly to user's creative requests. Combined with reliable speech

recognition, this makes it possible for users to describe the experience they want, much as the characters in *Star Trek: The Next Generation* interacted with the Holodeck.

Despite advances in SAR and immersive projection, existing systems remain limited by static, pre-authored content and require specialized technical expertise to create or adapt experiences. There is a clear need for tools that empower everyday users to author dynamic, personalized stories in their own physical environments without the barriers of head-mounted displays, complex setup, or professional design skills.

Our goal with *Storycaster* is to explore how **everyday users**, such as parents, educators, and casual storytellers, might author immersive room-scale stories through natural, conversational interactions. While professional creators have traditionally authored immersive experiences, *Storycaster* lowers the barrier for non-experts, to enable the creation of room-scale stories through intuitive, natural interaction. We envision a future where everyday environments (living rooms, classrooms, hospital rooms), become mediums for personalized creative expression, where anyone can generate these experiences on demand without technical expertise. To clarify user roles, we distinguish between (1) *Storyteller*: the participant(s) who speak to the system, issue commands, and shape the narrative, (2) the *AI Narrator*: which guides the story, generates dialogue, and responds to user input, acting as a co-storyteller, and (3) the *Audience*: Anyone present in the room who experiences the story; the storyteller and AI narrator are both part of the storytelling process, and the human storyteller can also be a part of the audience.

This user-centered framing led several key design choices. First, we chose to use voice as the primary interaction modality because it is intuitive and aligns with the natural, conversational nature of storytelling. Voice input allows users to shape their narrative (and any changes they make along the way) without relying on additional devices or screens. Second, *Storycaster* maintains users' awareness of their physical surroundings by augmenting the real environment, rather than replacing it. Unlike traditional virtual reality systems, which often immerse users in fully synthetic worlds, *Storycaster* enables users to remain present and engaged in their everyday environment, and have the ability, with a simple voice command, to change their everyday surroundings. Third, we target open-ended improvisational storytelling rather than pre-structured, choose your own adventure narratives. While the system does offer branching choices, its primary aim is not to guide users through a pre-determined narrative, but rather help them invent their own through changing settings, transforming objects, introducing characters, and redirecting the plot.

Previous immersive narrative experiences rely on pre-authored content. For example, the KidsRoom transformed a physical space into a story inspired by *Where the Wild Things Are* through carefully designed visuals and interactions with physical props (e.g., “now pretend the bed is a boat!”) [4]. While groundbreaking in narrative immersion and interaction, the experience followed an entirely predetermined storyline, guiding kids through scripted interactions rather than giving them freedom to change the storyline itself.

By leveraging generative AI, *Storycaster* enables content to be created on demand, allowing environments, characters, and narratives

to evolve in direct response to the audience. Through voice command, users can change environments, reimagine physical objects in the room, and choose the storyline they want to be immersed all in real-time (this includes the generation of ambient sounds, dialogue, and visual projections). In doing so, *Storycaster* creates a uniquely participatory platform, where users do not just consume a story but actively co-author it. A key question is whether generative AI techniques can deliver an experience as compelling as more predetermined and hand-authored experiences such as the KidsRoom.

Storycaster, through an orchestration of numerous GPT 4.1 driven components, designs a unique experience for every story and every room. It guides users in creating a three act narrative, using live camera feeds of the room to produce cylindrical images that are mapped back onto the walls and objects in the room. In particular, we make the following contributions:

- (1) **Room Scale Audio-Visual Generation:** We present an image generation pipeline that uses generative AI techniques to create projection mapping effects that alter the look of objects in the room or create entirely new scenes. The pipeline combines SDXL with Depth Control Net and a cylindrical LoRA, taking as input live depth images of the room, and a prompt to produce cylindrical images that are projected in the room. Figure 1 shows some example generations. Additionally we use Stability Audio 1.0 [42] to create matching ambient soundscapes, rendered by a Dolby spatial audio system to position audio dynamically throughout the room.
- (2) **Narrator Agent:** We developed a Narrator agent that guides participants throughout the experience, enabling users to co-create stories in the room. Through simple voice commands picked up by a microphone in the space, users can interact with the agent, crafting a storyline with clear actions, setting, and dialogue. The agent orchestrates the experience by generating appropriate visuals, ambient audio, and different character voices for every scene.
- (3) **Object Editing:** We demonstrate the ability to edit at an object-level in the room, giving users the granular control to virtually augment or transform physical objects in context of an evolving narrative. We investigate the use of Grounded SAM on cylindrical images to segment out objects in the room, coupled with SD 1.5 inpainting to generate virtual overlays and map them on corresponding objects in the room.
- (4) **Evaluation of *Storycaster*:** We conducted a study with thirteen participants and three pilot sessions. Each user was able to create one or more stories with *Storycaster* in a 20-25 minute experience. Through a pre and post experience questionnaire and a semi-structured interview, we evaluated *Storycaster* in its ability to immerse users in a story, as well as determine which part(s) of the experience contributed most to user's overall enjoyment. Our findings offer insights for future open-ended immersive storytelling experiences, and their broader applications.

2 Related Work

2.1 CAVE Systems

Cave Automatic Virtual Environment (CAVE) systems [10] are projection-based systems that are designed to immerse virtual reality in a physical space. From CRT projectors to multi-projector CAVE systems with commodity hardware [1] to the StarCAVE [13], they have evolved over the years, adapting through technological changes. Most recently, CAVE systems have extended beyond projectors, replacing projection screens with flat panel displays [16]. In all cases, they are effective in having immersive room-scale VR environments that blend together virtual imagery with the user's physical space. However, there has been limited work, to the authors' knowledge, on CAVE systems where users can use generative AI techniques to deliver custom narratives and visuals in their environments. This is what *Storycaster* aims to address.

2.2 Immersive Storytelling

Immersive storytelling departs from traditional, linear, single-medium narratives by inviting audiences into fragmented stories that unfold across multiple forms of media [37]. Within the Metaverse, a digital environment enabled by virtual reality (VR), storytelling takes shape through four distinct dimensions—agency, interactivity, immersion, and community—that distinguish it from conventional approaches [18]. These qualities foster personalization and a sense of co-authorship, as each user actively shapes their experience. VR and mixed reality further expand the possibilities of narrative by enabling audiences to inhabit diverse scenarios and perspectives in both real and fictional worlds [14]. Projects such as RealityMedia demonstrate how VR can function as a testbed for narrative experimentation, where spatial context, collections, and artifacts become central to story construction in ways unavailable to traditional media [28]. Similarly, ConnectVR explores audience-driven storytelling through a trigger–action framework, allowing characters and objects to adapt their behaviors in response to user interaction [8]. In the realm of audio, there has been work on creating personalized soundscapes and immersive audiobooks [12, 38]. Despite these advances, a common thread persists: immersive narratives still depend on physical devices—such as headsets or mobile phones—for participation. Little work has yet explored how to craft immersive stories that unfold directly within physical space, without the mediation of external devices.

2.3 Narrative Creation with Generative AI

Generative AI has recently become a powerful tool for storytelling [27], enabling audiences to engage in back-and-forth conversations with AI narrators [7]. Such personalized narratives not only foster creativity but can also help individuals reflect on and process their own experiences and emotions [15]. To support this, systems like the Universal Narrative Model [19] propose standards that leverage the strengths of generative AI while preserving the author's intent and creative vision.

A variety of frameworks have been developed to explore narrative generation, including approaches to branching storylines [23, 46, 47], different kinds of narratives including movie scripts [51],

and the design of novel storyworlds [29, 34]. Tools such as StoryForge [48] provide user-friendly interfaces that balance human authorship with AI collaboration, while projects like ID-8 [2] integrate visual and audio elements to create rich, multi-modal experiences. Beyond written narratives, AI has also been a transformative force in interactive games [24], extending from early text-based adventures like Zork [32] to collaborative role-playing systems such as PANGea [6] and 1001 Nights [43].

2.4 Uniting Physical and Virtual Worlds

Immersion lies on a spectrum—from augmented to virtual to mixed realities—where tangible interaction remains a central tenet [39]. Inspired by Ishii's Tangible Bits philosophy, the goal is seamless integration between digital content and physical experience [26]. This is seen in early work with toys and augmenting these into children's storytelling [22].

Several systems generate VR environments based on real-world physical structures. VRoamer creates VR scenes that dynamically align with unknown building layouts, enabling safe continuous walking [9]. DreamWalker substitutes real-world walking with virtual cityscapes, overlaying virtual assets onto physical paths [45]. Sra et al. [41] procedurally generates VR environments from reconstructed indoor spaces. All of these systems leverage physical geometry to inform environment substitution within VR, maintaining safety during movement and enabling large scale-exploration.

While *Storycaster* shares the idea of grounding virtual content in physical geometry, it diverges in several ways: (1) *Storycaster* focuses on storytelling with no headsets, where users remain visually aware of the room they are in, (2) VRoamer/DreamWalker replace the environment, while *Storycaster* enables users to author environments, characters, and narratives through voice commands, and (3) *Storycaster* includes more granular control including object-level editing to turn physical objects into narrative elements (an interaction not present in these VR systems).

Other recent AI+XR systems, such as LLMR [11], ImagineAR [33], DreamCodeVR [20], Dreamcrafter [44], and VRCopilot [49], have explored scene authoring and editing in extended reality, often via natural language interfaces. However, these typically require external screens or head-mounted displays, and do not offer the in-situ, room-scale immersion that *Storycaster* provides. While some storytelling systems are beginning to integrate more object-centered interactions in order to shape narratives in physical space [3], these still lack the 360° immersion that a room provides. Work such as IllumiRoom [31] takes a step in this direction, creating immersive experiences for gaming on TVs, turning the surroundings of the device into the environment of the game itself. Similarly, Room Alive provides a framework for turning any room into an immersive gaming experience [30].

Storycaster is a novel step beyond these: it enables users to construct narrative worlds using simple voice commands, while spatial reasoning grounds the experience directly within the physical space. In doing so, *Storycaster* bridges the immersive room-scale presence of CAVE systems, the narrative flexibility and co-creative potential of modern immersive storytelling, the generative power of AI-driven narrative construction, and the seamless physical and digital world integration.

By combining these threads, *Storycaster* seeks to redefine immersive storytelling as a physically embodied dynamically controllable narrative experience, rooted in the user’s real environment itself.

3 Storycaster

Storycaster brings together various integrated components to create an immersive storytelling experience. The system includes: (1) an image generation pipeline that generates 360° cylindrical images from live camera feeds and maps them back on the walls and objects in the room, (2) a Narrator Agent that guides participants throughout the experience, and (3) object-editing capabilities that allow users to transform physical objects into virtual overlays contextualized within an overarching story. The overall architecture of the system is described in Figure 2, and the prompts used by each module are discussed in the Appendix.

When first entering the room, users are welcomed into a projected library environment, accompanied with gentle background music. Bookshelves line the walls, encouraging users to “choose” a book from the shelf or make a story of their own. They are then guided through a tutorial section that demonstrates the room’s capabilities: image generation, ambient audio generation, and object-level editing. Participants interact by speaking aloud; their input is captured by a room microphone, transcribed with Azure Speech to Text services, and processed by the Narrator agent. A chime sound signals after a user’s input has been registered, reassuring them that their voice has been heard. Additionally, users are asked if they have brought in a personal object; this object is later integrated into the stories they create, acting as a “Macguffin” in the story, a central talisman for which the story revolves around [36]. Similar work was done in “Rosebud” that treated a child’s toy as a central object in a story they created [21].

Once participants are comfortable with the system, they transition into the actual story experience. The experience follows a three-act structure commonly used in film: The Setup, The Confrontation, and the Resolution [17]. Users begin by engaging in a conversation with the Narrator that helps them craft the initial premise for their story, the agent probing for characters, the setting, and the main action taking place to situate the story. Once the narrator has enough information to create the first act, it asks the user for confirmation, and proceeds, generating the script, visuals, ambient audio, and the different voices for characters in the scene.

At the end of each act, users are prompted to either continue the story or modify elements in the room. If they choose to continue the story, the Narrator presents three possible directions for the narrative, while also allowing users to provide their own storyline or let the Narrator decide. This repeats for all three acts of the story, culminating in users creating a new story or ending the session.

3.1 Room Features

For sensing and display there are six projectors and four Azure Kinect cameras mounted in the ceiling of the room. The cameras take in both color and depth images. The cameras and projectors are mounted to roughly cover the full room below the ceiling and are calibrated as in [30]. Additionally, there is a Dolby Atmos spatial audio system set up for sound playback as well as a microphone for user input.

The current implementation was installed in a single-calibrated projection-mapped room; however, the underlying approach is designed to be adaptive. *Storycaster* can theoretically operate in any room equipped with calibrated projectors and depth cameras, since the pipeline does not rely on room-specific characteristics or dimensions. However, environmental constraints such as cluttered spaces, or non-rectangular rooms can reduce projection quality. Future work will explore a more portable lightweight version of our setup.

3.1.1 Image Generation. The projectors display imagery that can subtly alter the look of the room or more dramatically transport the viewer to an altogether different space. These effects are created to follow and support the narrative, using generative AI techniques such as stable diffusion.

To generate the images, a cylindrical (360°equirectangular projection) depth image of the room is first assembled from the multiple live depth camera images. We opted for cylindrical images as this allows for a one-shot approach of generating an image that can be projected throughout the entire room, rather than generating multiple separate views. A depth ControlNet uses the cylindrical depth image as guidance to encourage image generation to follow the physical structure of the room, including the walls and furniture.

We utilize two models for image generation: SDXL for generating full image scenes and SD 1.5 for object-level inpainting. In both cases, the live cylindrical depth image of the room and the prompt is passed into the model which is augmented with a depth ControlNet [50], and LoRA tuned to generate cylindrical images. Through experimentation, we used a value of 0.76 for the depth ControlNet strength value for SDXL, and 0.54 for inpainting with SD 1.5 (refer to Appendix 14 for these results). This produces 360° cylindrical images that are then upscaled and projected back onto the four walls of the room, using projection mapping techniques. The image generation process is illustrated in Figure 3.

While the depth control net ensures that wall surfaces and room geometry influences image placement and perspective, the semantic meaning of physical objects is not explicitly modeled during scene-wide image synthesis, unless explicitly invoking the object-editing interaction. The prompts rather focus on describing the setting of the scene (“deep forest at dusk”, “library”) and letting the model take creative liberties with what each object becomes. This design choice allowed generations to remain robust across diverse stories but also resulted in mixed user feedback regarding object-level coherence. We discuss this further in Section 6. When users initiate object-level editing (e.g., “change the table into a black hole”), the system incorporates the semantic label of the physical object into the inpainting prompt, enabling targeted transformation. This approach allows users to personalize their stories by re-purposing familiar items as narrative anchors. Our user study found that object-level editing enhanced personalization and narrative engagement, as participants enjoyed transforming real objects into story elements (see Section 5).

3.1.2 Audio. For generating ambient audio, we utilize Stability Audio 1.0, with outputs designed to loop continuously for seamless playback.

Additionally, we define an AudioPlayback API that manages multiple sound sources in the room, each defined with a unique

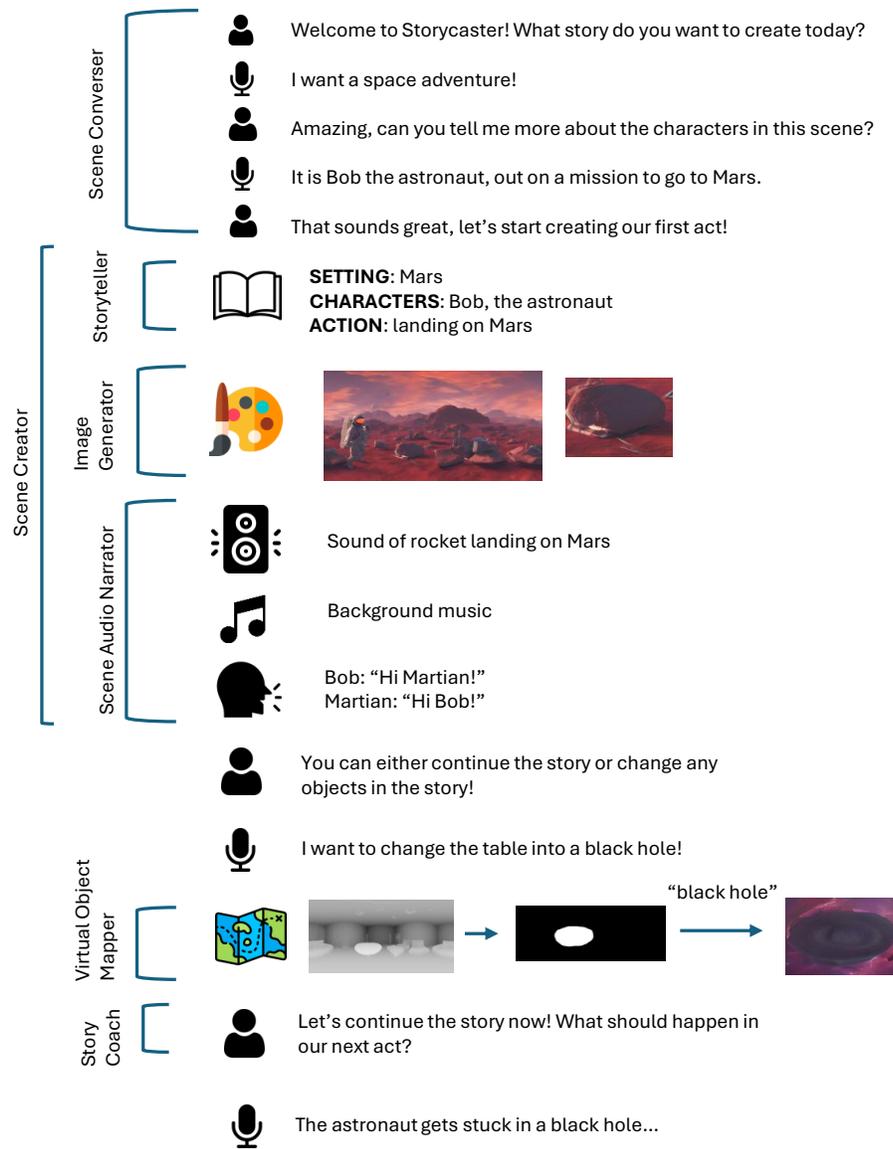


Figure 2: System overview of *Storycaster*. The Narrator orchestrates the storytelling process, prompting the user for input and guiding them through three acts. The Scene Converter gathers initial story elements, which are passed to the Scene Creator. Within the Scene Creator, the Storyteller generates narrative text, the Image Generator creates visuals, the Virtual Object Mapper overlays virtual counterparts onto physical objects via inpainting, and the Scene Audio Narrator produces ambient audio and character dialogue. The Story Coach manages transitions between acts, allowing users to either continue the narrative or change objects in the room. Segmentation, inpainting, ambient audio generation, and image generation all happen on the server side. This process repeats for three acts, after which the user can choose to begin another story or conclude the session.

name, a volume, and spatial coordinates. The API allows multiple sources to be added, replayed, or removed. Each sound source is instantiated as a separate connection, and played sequentially through the spatial audio system.

3.1.3 MCP Server. In order to seamlessly integrate the LLM driven storytelling components with the room, we implement a

Model Context Protocol (MCP) architecture for our server-client setup. The server exposes several APIs for image and audio generation, allowing the client to orchestrate story structure while offloading computation-intensive tasks to the server.

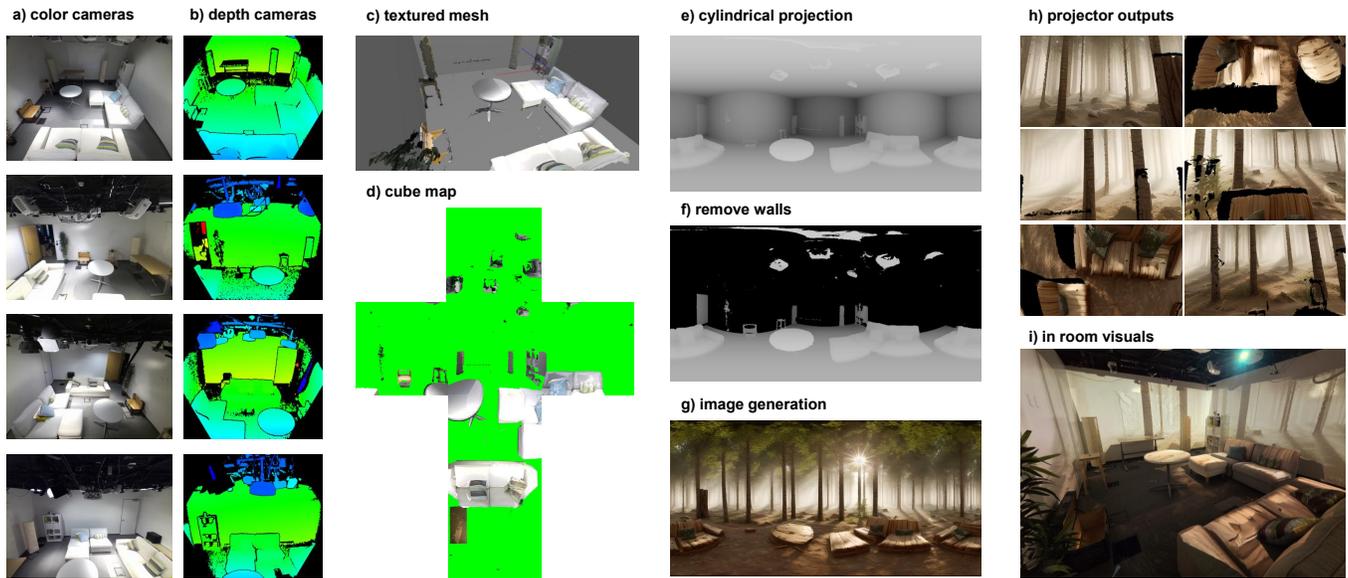


Figure 3: *Storycaster* combines multiple color and depth cameras to generate visuals that follow the physical features of the room. Calibrated color (a) and depth cameras (b) (rendered as false color) are combined to create a textured mesh that serves as a realtime geometric model of the room (c). This model is rendered to a single cube map (d). The cube map z-buffer is re-rendered as a cylindrical (or 360°) depth image (e). To render an exterior or outdoor scene, walls may be optionally removed (f). Stable diffusion is applied with a depth ControlNet conditioned on (f) and a 360° LoRA model. ControlNet effects are masked (eliminated) for pixels that are zero in (f), allowing for a more natural horizon effect in exterior (outdoor) settings. The generated image (g) is upscaled 4x (not shown). The image displayed at each projector is calculated from the geometric model of the room, using intrinsic and extrinsic parameters of the projectors (h). Black pixels correspond to regions of the geometry that are not visible in (e) (i.e., they are occluded by another object). The final projected visuals create 360° views into an exterior virtual space and change the appearance of objects in the room (i).

3.2 The Narrator Agent

The Narrator agent functions as a main orchestrator, managing multiple LLM-driven (GPT 4.1) components. These include a Scene Converser, and multiple instances of a Scene Creator, each consisting of a Storyteller, Image Generator, Scene Audio Narrator, a Virtual Object Mapper, and a Story Coach. The overall architecture and module interaction are illustrated in Figure 2.

The Narrator includes a speak function that is called for each piece of text generated. This function takes in the text to be recited, the type of voice, and mode, which varies based on whether the system is in tutorial or storytelling mode. Each mode utilizes a different style prompt: the tutorial mode uses a more relaxing calm voice, while storytelling mode uses an engaging, friendly delivery. This style prompt and the text is passed into GPT’s text to speech model (TTS), which outputs audio at 24 kHz. The audio is resampled to 48 kHz, and sent to the AudioPlayback API to be played through the room’s speakers. We chose to use the TTS model instead of the GPT Realtime model, despite its more natural conversation flow, due to the level of control and structure needed for the experience.

When the Narrator prompts users for input, they can simply speak in the room. A microphone captures their voice, which is streamed via a websocket URL to Azure Speech-to-Text services for transcription. To handle pauses or hesitations without cutting off

transcriptions prematurely, we increase the end-of-segmentation silence timeout to 3 seconds. Once the speech recognizer determines the user has finished speaking, it returns the final transcription, closes the stream, and plays a chime sound to signal the Narrator has successfully received the input.

We recorded performance metrics to characterize system responsiveness. It took approximately 7.3 seconds to produce images (both scene-level and object-level) including upscaling, and 12 seconds for ambient audio generation. During the image generation, users were able to see the images form in the room from the initial model generation to the eventual final version after upscaling. In terms of generations, during our user study, all generations were generally successful and correspondent to user’s requests determined by their responses to an interview (see Section 5 for details). Future work includes reducing these latency values so the experience is as seamless and smooth as possible.

Each module of the Narrator agent is described in detail in the following sections:

3.2.1 Scene Converser. When first starting the story creation process, the Scene Converser asks the user what kind of story they would like to create. It engages in a back and forth conversation with the user to craft an initial premise for their tale, probing for

four key elements: (1) the setting, (2) characters, (3) main action, and (4) the mood and tone of the scene. The Converser is designed to keep the instruction lightweight, avoiding excessive questioning as details can be refined later in the story.

After each user response, the Scene Converser evaluates whether it has enough information to proceed creating the scene. This is determined by both the information the user provides, as well as the number of conversational turns. If yes, then it will respond with “READY”; otherwise, it responds with “NEED MORE”. Once “READY” is returned, the Scene Converser crafts a personalized prompt asking the user if they are ready to proceed with bringing their story to life. It then analyzes the user’s response to identify if they have agreed or want to continue the story (such as “yes”, “I am ready”, “let’s go”). If it determines positive confirmation, it creates an instance of a Scene Creator and starts the scene creation process. If the user wishes to add more details, the Converser continues conversing until the user is satisfied. For users who already know exactly what story they want to tell, they can skip the back-and-forth with the Converser entirely by indicating readiness, after which the Converser confirms, and proceeds creating the scene.

3.2.2 Scene Creator. Once a user is ready to create a scene, an instance of a Scene Creator is created. The Scene Creator is responsible for three components: the Storyteller, in charge of the story text; the Image Generator, in charge of image generation and object editing capabilities; and the Scene Audio Narrator, in charge of generating both ambient soundscapes and character voices. Together, these classes transform user’s input into a multimodal storytelling experience.

3.2.3 Storyteller. The Storyteller is in charge of generating the actual story script. It takes in the following inputs: (1) the **current act** (first, second, or third), (2) any **previous scenes** generated, (3) the **original storyline** from the user’s initial interaction with the Scene Converser, (4) the user’s **personal object** description from the tutorial, and (5) the current **user input**. The original storyline and the current user input are the same if creating the first act.

Using these inputs, a detailed system prompt is crafted to generate a scene based on these personal preferences. In order to make the stories feel more like a play with heavy dialogue, we use few-shot prompting of various movie scripts from the IMSDB movie script database [25]. We hand-selected four such scripts from this database, and extracted out the first 25 lines of two random scripts for each scene generation, ensuring variety across story structures. In particular, we aimed to emphasize setting, dialogue, and actions in the generated stories—something that movie scripts do best. Guided by these examples and detailed instructions on key storytelling principles—such as prioritizing character dialogue, using concise visual descriptions, and revealing backstory and motivations behind character’s actions—the Storyteller generates a 8-10 sentence scene.

3.2.4 Image Generator. Once the Storyteller generates the current scene text, it is passed to the Image Generator, which is responsible for both the wall projections and object-level generations in the room.

For room-scale visuals, the scene text is used to generate an appropriate prompt for SDXL, the image generation model. The

system prompt for this conversion was carefully crafted, and includes few shot examples pairing different story excerpts with corresponding visual prompts of no more than 50 words. Additionally, based on insights from pilot studies, we added explicit instructions to favor “saturated bright colors” as we observed this improved the vibrancy of the generated images in the room.

The generated prompt is also used to classify if there should be a removal of walls in the generation. This was determined by classifying the specified scene as outdoors or indoors; outdoor environments would enable the removal of walls and indoor environments would not. Enabling the removal of walls for outdoor environments allowed for visuals to extend beyond the four walls and create a sense of openness. Indoor environments, by contrast, preserve wall constraints to maintain spatial consistency.

Finally, the refined prompt, classification output, and a white mask (ensuring the full depth image of the room is used) are passed into the MCP image tool to generate the final projection.

3.2.5 Virtual Object Mapper. The Virtual Object Mapper, a component of the Image Generator, is responsible for generating appropriate virtual object mappings to the physical objects in the room that align with the current story. This includes detecting objects, generating mappings between physical and virtual items, and producing inpainting prompts to pass into the server-defined inpainting tool. This process is performed for every scene, and users may also choose to remap objects of their choice between acts.

For object detection, we found that Grounded SAM did not consistently produce reliable masks (please see Section 3.3). Instead, we opted for a more manual approach where we hand-crafted masks of various objects in the room (e.g., table, ottoman, lamp, and sofas). Each mask file is labeled in the format `maskobjectwhite.png`, enabling straightforward retrieval.

Once objects are identified, the system generates mappings using the following inputs: (1) The current scene, (2) the previous scene for consistency between two scenes, (3) the detected physical objects in the room, and (4) the previous mappings generated (if applicable). Using few shot prompting of different physical objects and potential virtual overlays they can be converted to, the system prompt lays out mapping guidelines that emphasize story consistency, physical compatibility with the shape of the object, and continuity between previous mappings. The output includes both a mapping of physical to virtual objects, and an optimized inpainting prompt.

Finally, for each object, its corresponding mask is retrieved and converted to a base64 string. This string, combined with the inpainting prompt, is passed into the image tool to generate overlays, effectively transforming real objects in the room into story-appropriate virtual artifacts.

3.2.6 Scene Audio Narrator. The final component of a Scene Creator instance is the Scene Audio Narrator, in charge of both the ambient audio generation and the voice acting for a scene.

To generate ambient audio, the scene narrator takes in the generated scene text from the Storyteller and using a carefully designed system prompt with few-shot examples of story-audio pairs, constructs an appropriate concise description of the desired soundscape. This prompt is then passed into the server-side ambient audio tool to generate the sounds.

Additionally, background music is played in between acts through searching Spotify [40] for top 5 playlists that meet the prompt “background instrumental music” and plays a song from the top search. This way, users do not hear the same music for each act, and in between stories.

Once the ambient sound has been created, the Scene Audio Narrator produces dialogue for the various characters in the scene, using the GPT TTS model. This process consists of (1) parsing the scene text into individual lines, each one assigned to a specific character (or the Narrator), (2) assigning a distinct vocal style to each character based on dialogue content and story context (passed in as the “instructions” parameter in TTS), (3) assigning a voice for each character from GPT’s voice selections (alloy, echo, fable, onyx, nova, and shimmer), determined on the scene text and their voice style assignment, and (4) rendering the recitation of voices. Additionally, to reinforce spatial immersion, each character is assigned a fixed location in the room, depending on the number of the characters in the scene. This allows for conversations to play out across the room and for characters to speak from different locations. Consecutive lines by the same character are grouped together to streamline recitation.

Dialogue playback is managed by the server’s AudioPlayback API, which takes as input the sound source name, volume, and spatial position. All TTS instances for a scene are generated simultaneously, ensuring synchronized timing, and played sequentially to avoid latency between lines. Once narration concludes, all connections are closed, finalizing the scene’s audio performance.

3.2.7 Story Coach. After a scene concludes, the Narrator prompts the user to either continue the story or modify any objects in the room. If a user indicates they want to continue the story, the story transitions to the next act and the Narrator calls the Story Coach. If the user indicates they want to change an object in the room (e.g., “I want to change the ottoman to a campfire”), their request is parsed to identify the physical object and its desired virtual counterpart. The corresponding mask is retrieved, and the inpainting API is directly called with these inputs.

For continuing the story, the Story Coach takes in the following inputs: (1) the original storyline established with the user’s initial interactions with the Scene Converser, (2) how many acts have happened so far and what act we are currently at, and (3) the most recent scene text. From these, the Story Coach provides three possible directions for the narrative, tailored to the current act. Users can then: (a) select one of the provided options, (b) propose their own idea, or (c) delegate the decision to the Narrator. For option (a), an additional LLM call extracts the full text of the selected option, which is then used to generate the next act. For option (b), the user’s idea is passed directly as input with no further questioning. For option (c), the Narrator generates the next part of the story, building on prior events.

3.2.8 Wrapping Up. The described process above repeats until all three acts are complete. After the third act, the Narrator congratulates users on completing their story, and asks if they would like to create another one. If yes, a new instance of the Narrator is launched; otherwise, the experience ends.

3.3 Object Editing

One of the key interactions in this experience is the ability for users to edit at an object level. This granular control allows users to select a physical object in the real world and turn it into one in the story world, while still preserving the object’s underlying shape and size. See Figure 4 for examples taken from our user study.

To detect physical objects, we apply Grounded SAM to segment out objects and their masks from a cylindrical image of the room. By using the label “all objects,” the model extracts masks for every visible object without requiring predefined categories. These masks, along with the room image, are then passed to GPT-4o to generate textual descriptions of the detected objects.

In practice, we find that grounded SAM performs well for objects in the central region of cylindrical images, and less so for objects on the top or bottom. For items of similar shapes and sizes (ex. ottoman and table), we found that Grounded SAM does not return consistent results, oftentimes confusing one for the other. For this reason, we hand-crafted masks for key objects in the experience so we could be sure that the object a user requests is the one being transformed.

Interestingly, we also discovered that grounded SAM can be used to identify *virtual* objects that are not in the room itself. By passing in the name of a virtual object and lowering the similarity threshold (we used 0.1), the model correctly identifies objects in the room that look most closely to the specified one. For instance, as seen in Figure 5, prompting “boat” returned the mask of a sofa, while prompting “vine” highlighted the ladder. This capability enables flexible mappings between real and imagined worlds, and to our knowledge, is a novel contribution to the system’s semantic reasoning and interpretive flexibility, particularly in the context of object-level editing.

Object-level editing was included not only to demonstrate technical feasibility but to explore a new feature to support narrative making. Unlike scene-wide visual changes, object transformations allow users to re-purpose familiar physical items (sofas, tables, lamps) as narrative anchors. Even when generative quality varied, object editing created a connection between the virtual and physical worlds, enabling interactions such as turning an ottoman into a campfire, or a sofa into a boat. This form of re-contextualization is rare in prior AI + XR systems, and is a key distinction of *Storycaster*.

4 User Study

We conducted a user study to explore how users felt about the experience and the type of content they would generate. In particular, we were interested in exploring (1) how immersed users felt in the experience, (2) what specific features of the experience added the most to their overall enjoyment, and (3) how users reacted to the physicality of the room while still being able to create a virtual world. Figure 6 shows examples of the outputs generated by participants during the study, which we further describe in Section 5.

This evaluation was intentionally framed as a usability exploration, rather than a comprehensive benchmark against screen-based storytelling or alternate XR interfaces. Our goal was to understand how users interacted with the system, and what new forms of immersion a system like this enabled. We acknowledge



Figure 4: Examples of object editing from user studies. Participants selected real-world objects (table and ottoman, shown with their masks on the left) and reimagined them as virtual objects within their stories. The resulting generations (right) illustrate diverse mappings such as a campfire, quilt, black hole, burrito, and flowers, demonstrating the flexibility of object-level transformations in *Storycaster*.

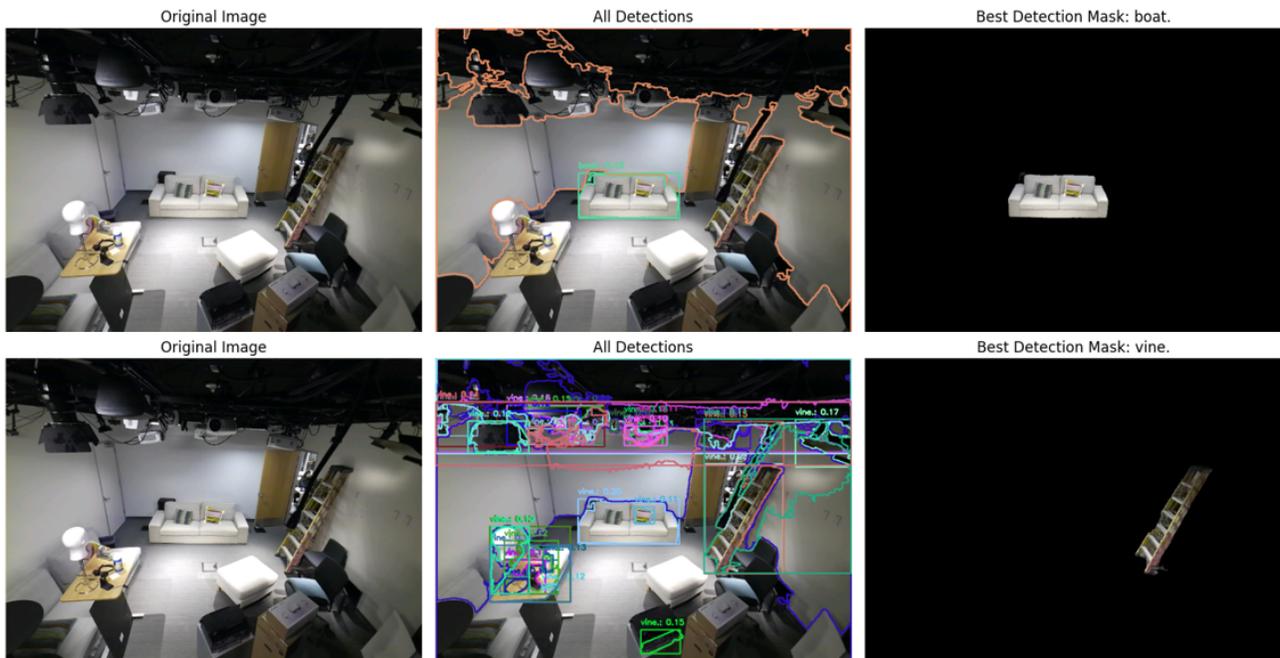


Figure 5: Examples of using Grounded SAM to identify virtual objects by passing in object prompts with a lowered similarity threshold (0.1). Top row: prompting “boat” highlights the sofa as the closest visual match. Bottom row: prompting “vine” highlights the ladder. From left to right: original room image, all detected masks, and the best detection mask. This demonstrates the ability to flexibly map real objects to virtual counterparts, enabling novel object-level editing in immersive storytelling.

the value of comparative baselines, such as 2D generative storytelling, scripted versus generative conditions, and with/without object-level editing. These comparative conditions are important for quantifying *Storycaster*'s unique contributions and will be the focus of future studies. In this initial work, we prioritized open-ended, in-situ interaction to capture authentic user experiences and preliminary behaviors.

4.1 Participants

A total of 13 participants participated in the study, with three additional one being pilot studies. Age was collected in ranges (e.g., 25–34, 35–44). Most participants were in the 25-43 age group. To estimate the average age, the midpoint of each range was used. The resulting average age of participants was approximately 34.1 years. Out of the 13 participants, there were 5 females, 7 males, and one gender variant/non-conforming participant.



P31: Historians in East Germany stumble into something unexpected.



P32: (N) Historians discover a device from an old station.



P33: Historians uncover a hidden Cold War-era room beneath the station.



P51: Story from Ramayan where Ram, Laxman, and Sita are in the forest; they see a deer and Sita insists she wants it.



P52: (N) Ram goes deep into the forest to find the deer for Sita.



P53: The deer transforms, revealing a hidden threat to Sita.



P61: Participant preparing dinner with her parents and cat in their kitchen



P62: The cat and dog team up for a wild tomato chase.



P63: Mom declares a truce and serves garlic bread to everyone.



P71: A researcher returning to Earth after discovering we are all alone in space



P72: (N) Researcher returns home after a lonely mission.



P73: (N) Researcher drinks a cup of coffee, bitter about his journey.



P81: A group of friends explore a moonlit castle.



P82: The group of friends discover a cave that has a dragon.



P83: The friends help the dragon solve a problem in the cave.



P91: User's mentors are professors at Hogwarts with him and three of his fellow interns as students, and two other students in a romantic relationship.



P92: One of the students throws a banana at the couple.



P93: The couple reveals a secret about their magic.



P101: 3 snow-cone makers live in a castle made of ice that is affected by oncoming darkness.



P102: The snow cone makers try to see the king and are barred out by the guards.



P103: (N) The king lets the snow cone makers enter the gate.



P111: The user and his sea lion nemesis compete to open a burrito store and a coffee shop in San Francisco.



P112: The sea lion rats the user out to the safety inspector, but then the health department shuts them both down.



P113: (N) The sea lion and the user team up together to open a burrito-coffee business.



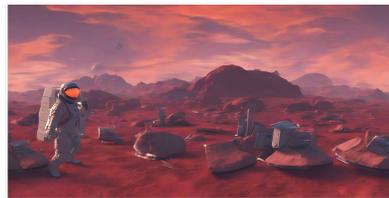
P121: The user's parents send her off on a spaceship with her boyfriend.



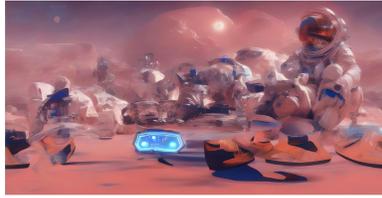
P122: (N) The user says goodbye to her parents as she boards the spaceship with her boyfriend.



P123: (N) The user and her boyfriend start up the ship and go on their space quest.



P131: The user is exploring Mars.



P132: A sudden Marsquake shakes the ground, revealing a hidden hatch.



P133: (N) The user finds strange symbols on the hatch, and starts exploring.



P141: The user is a bird soaring with other birds exploring the views of the Dolomites and interacting in acrobatic challenges.



P142: The birds stumble upon a mysterious, swirling cloud tunnel.



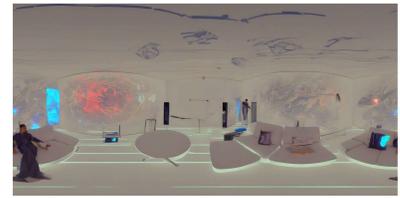
P143: (N) The birds fly through the tunnel, laughing as sunlight flickers.



P151: A chopper saves a Jedi from Star Wars.



P152: (N) The two protagonists Lira and Jax escape drones, holding a token.



P153: Lira reveals a hidden Jedi signal, guiding them to safety.



P161: Three people are stranded on a shipwreck, playing rock paper scissors to see who would be the next meal.



P162: A water spout drops a crate of mysterious canned food aboard, the food making them hallucinate.



P163: Their hallucinations become shared and they all start living an alternate life where they are gnomes.

Figure 6: Generations from user studies ($n = 13$). Each row represents one participant, with three images corresponding to the acts of their story. Labels indicate participant ID and act number (e.g., P51 = Participant 5, Act 1). Images marked with (N) denote acts where the user allowed the Narrator to decide the direction of the story. The examples highlight the diversity of user-authored narratives, ranging from historical explorations and mythological tales to domestic scenes, space adventures, and fantasy worlds, demonstrating the flexibility of *Storycaster* in supporting a wide variety of storytelling themes.

All subjects reported interacting with AI systems on a daily level for technical tasks (coding, writing, data analysis), creative tasks (brainstorming, prototyping, generating images/video), and support tasks (summarization, text correction, accessibility, education). Most participants engage with stories daily or weekly, with a few reporting less frequent or specialized engagement (e.g., only occasionally, rarely, or through specific media like video or podcasts).

All subjects gave written informed consent prior to participating (reviewed and approved by the Institutional Review Board), and were given a 25 dollar gift card. Participants were recruited through various mailing lists.

4.2 Experimental Design and Procedure

We conducted an in-person experiment in the room. Users would first come in, and fill out a pre-survey prior to the experience asking basic demographic questions, their usage of AI systems, and how often they engage with stories. After filling this out, users participated in a 20-25 minute experience interacting with *Storycaster*. One door of the room was left open, so a staff member could listen in on what was playing and being said in the room, and monitor the generations. After the experience was over, the investigator came back and asked the user to fill out a post-experience questionnaire. This questionnaire asked them to rate their experience on a Likert scale from *Strongly agree* to *Strongly disagree* for multiple aspects of the experience, including visuals, audio, and story arc. After the user was done filling out the questionnaire, we conducted a semi-structured interview asking users about the overall experience. This included questions about their immersion, their interaction with the narrator agent, editing objects in the room, and applications they see a system like this being used for. Participants were encouraged to give feedback and suggestions at the end of the interview. This entire process took approximately 50 minutes to 1 hour.

4.3 Analysis

We used a mixed-methods approach, integrating quantitative insights from the questionnaire responses and qualitative feedback from semi-structured interviews and open-ended questionnaire items. We conducted a collaborative thematic analysis of the qualitative data, following Braun and Clarke's [5] six-phase process. Our team read and coded all interview transcripts and open-ended responses, identifying key aspects such as immersion, personalization, narrative interactivity, visual clarity, story structure, and feedback/agency. Codes were grouped into themes through iterative discussion and refinement using a shared Figma board (see Appendix Table 3 and Supplementary Materials). We defined each theme with supporting quotes and integrated them into the results, connecting qualitative insights to quantitative findings and informing our design recommendations described in Section 6.

5 Results

5.1 Quantitative Findings

5.1.1 Perceived Impact of Generative Aspects on Experience. To assess which components most enhanced participants' experiences, we collected ratings across five dimensions: 1) visual generation, 2) audio generation, 3) narrative generation, 4) object-level generation and 5) narrator delivery. Each statement followed the format:

"The [aspect] of the room setup enhanced my experience the most." Responses were collected using a 5-point Likert scale.

As shown in Figure 7, the Narrator was the most positively rated component, with strong consensus. Visuals also showed a strong positive trend but elicited the most polarized opinions since a small subset of participants expressed disagreement. Audio demonstrated consistent positive-to-neutral reception and in contrast, Object received the most mixed feedback with both positive and several negative responses. Narrative leaned positive but had a several neutral responses. These trends are further supported by the descriptive statistics in Table 1.

In addition to these metrics, participants also provided ratings for the smoothness of the transitions between different visuals. Compared to the five dimensions, participants' answers to the visual transitions were less positively received (total positive was 23.1%). The most frequent response (Mode (\hat{x})), was *Neutral*, and the median response (\hat{x}) was 3.0. The range of responses was 3.0, suggesting high variability in participant sentiment, which might have contributed to polarized "visual" dimension results.

5.1.2 Enjoyment Ratings Across Dimensions. We also analyzed enjoyment ratings to complement the dimension that enhanced the experience the most. As shown in Figure 8 and Table 2, enjoyment ratings largely mirrored the enhancement findings except for the "Object" dimension, which showed more positive trends. We subdivided this dimension in two categories: 1) personal object category and 2) transformation of objects in the room. More specifically, the questions asked were: *"I enjoyed seeing my object become a part of the story."* and *"I enjoyed transforming objects in the environment"*, which showed a more favorable distribution than *"The object-level generation aspect of the room setup enhanced my experience the most."* In addition, we also asked participants to rate how much their object *"played a role in the story unfolding"*, which received the same total satisfaction at 62.5%.

5.1.3 Immersion, Ease of Use and Future Usage. In addition to evaluating the 5 dimensions of the system, we also conducted a quantitative analysis of user responses across three key metrics: 1) Immersion, 2) Ease of Use and 3) Adoption intent. The results (shown in Figure 9), indicate that the experience was positively received, with both immersion and ease of use achieving 84.7% positive responses. All participants except one agreed or strongly agreed they would use such a system for their own use.

5.2 Qualitative Findings

This section describes the eleven sub-themes that emerged from our thematic analysis, categorized in three big themes: 1) Core System Feedback, 2) General Feedback and 3) Future Design Systems. Each theme is supported by quotes and integrated interpretations, highlighting how participants engaged with the system and how these insights complement our quantitative findings earlier described.

5.2.1 Core System Feedback. Six sub-themes emerged from the overall system feedback: Visuals, Audio, Narrator, Story, Object Editing, Personal Object Integration. These themes align with patterns observed in the quantitative analysis (Table 2) and are described below:

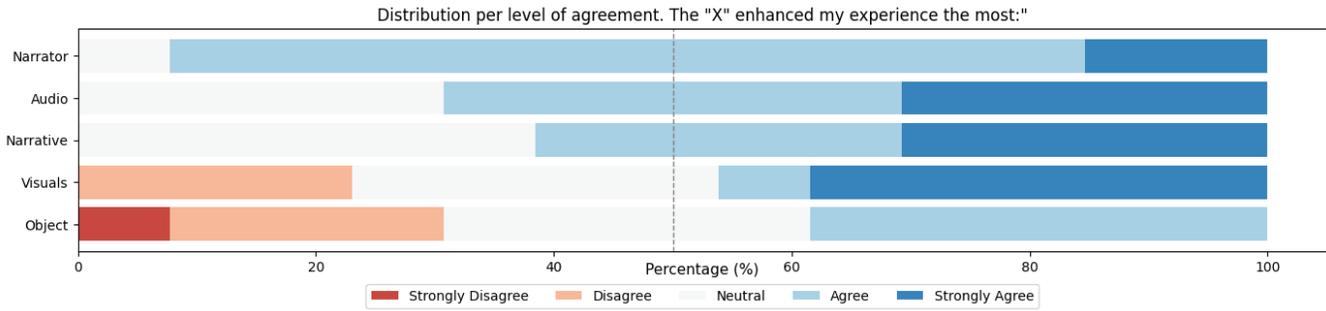


Figure 7: The stacked horizontal bar chart displays the distribution of Likert scale responses for five experience dimensions, showing which aspects participants felt enhanced their experience the most. The color scheme uses red tones for negative responses (Strongly Disagree, Disagree), neutral gray for Neutral responses, and blue tones for positive responses (Agree, Strongly Agree). The 50% reference line separates negative from positive sentiment areas.

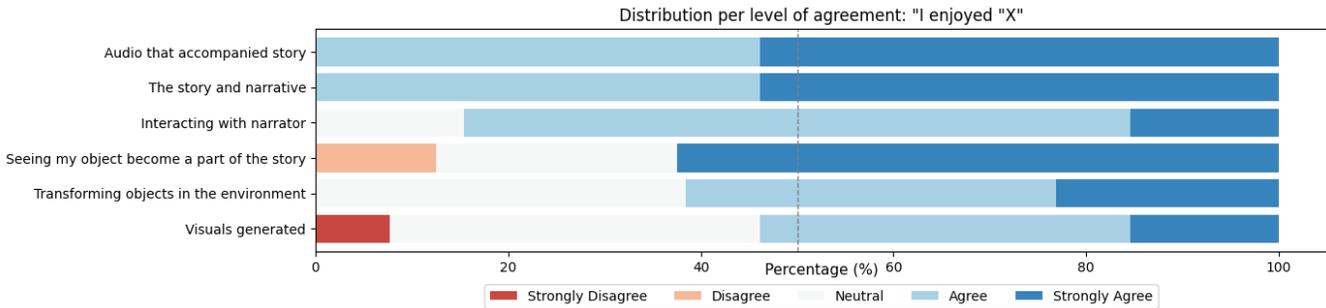


Figure 8: The stacked horizontal bar chart displays the distribution of Likert scale responses for how much participants enjoyed different aspects of their experience. The color scheme follows the same format as Figure 7.

Table 1: Descriptive statistics for perceived experience enhancement across dimensions.

Dimension	Mode (\hat{x})	Median (\tilde{x})	Range	Skewness (γ_1)	Positive Responses
Narrator	Agree	4.0	2	0.262	92.3% (76.9% Agree, 15.4% Strongly Agree)
Audio	Agree	4.0	2	0.000	69.3% (38.5% Agree, 30.8% Strongly Agree)
Narrative	Neutral	4.0	2	0.164	61.6% (30.8% Agree, 30.8% Strongly Agree)
Visuals	Strongly Agree	3.0	3	-0.012	46.2% (7.7% Agree, 38.5% Strongly Agree)
Object	Agree	3.0	3	-0.591	38.5% Agree

Table 2: Descriptive statistics for perceived enjoyment across dimensions.

Dimension	Mode (\hat{x})	Median (\tilde{x})	Range	Skewness (γ_1)	Positive Responses
Audio	Strongly Agree	5.0	1	-0.175	100% (46.2% Agree, 53.8% Strongly Agree)
Narrative	Strongly Agree	5.0	1	-0.175	100% (46.2% Agree, 53.8% Strongly Agree)
Narrator	Agree	4.0	2	0.000	84.6% (69.2% Agree, 15.4% Strongly Agree)
Personal Object	Strongly Agree	5.0	3	-0.895	62.5% Strongly Agree
Object-Transform	Agree	4.0	2	0.307	61.6% (38.5% Agree, 23.1% Strongly Agree)
Visuals	Agree	4.0	4	-0.885	53.9% (38.5% Agree, 15.4% Strongly Agree)

Visuals: For visuals, participants said that when good, they helped place themselves in the story. However, the visuals were

oftentimes vague with either unclear content, blurriness, or low-resolution images. This data is also reflected in the polarized results for the "Visual" dimension reported in the quantitative section. For

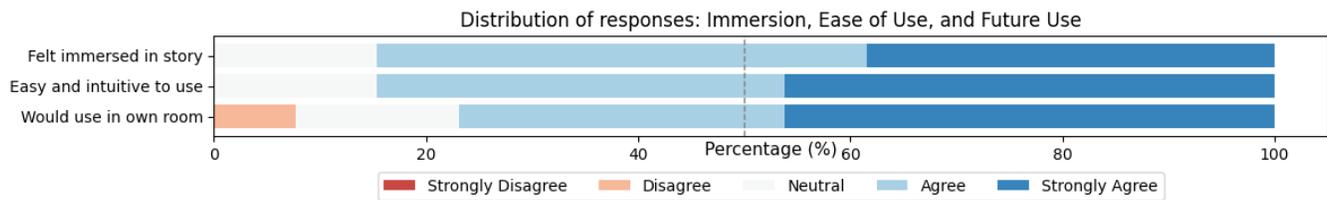


Figure 9: The stacked horizontal bar chart follows the same color scheme as Figure 7 and displays the distribution of Likert scale responses for immersion, ease of use, and adoption of the room, showing how participants rated these critical usability and environmental aspects of their experience.

instance, P5 states that when the visuals were good, “it really helped me just place myself in the story”, but that “sometimes it’s not clear, like, what is it that they’re trying to show”. While some disliked the blurriness of some images (P11), others liked that it made it easier for them to imagine (P12). Additionally, participants had mixed reactions on the abstract nature of some of the generated images. While some participants wanted images to be more photorealistic (P9), others found the abstractness of them to be conducive to imagination. P12 compared them to comics, stating that comics are also intentionally abstract and allow people to fit the image to their own personal scenarios. A couple of participants (P10,P12) were pleasantly surprised by how artistic the outcome visuals were. Although the visuals did not match what they had initially imagined, they still found them satisfying and acceptable.

Object Editing: Similar trends were seen for object-editing. While participants liked the granularity of the control and segmentation, the poor quality of the object visuals limited its overall impact. For instance, P10 said “the objects I often didn’t know what they were, but they melded in in conception”. Similarly, P10 noted how funny it was that she “said feast and it heard beast and this became covered in fur. So generally, I think I didn’t always know what the objects were supposed to be.” Others stated that while the object-editing capability was easy to interact with, they did not feel like it added much to the story. P6 mentioned that it did not work too well but could imagine object editing being the most impactful feature. We imagine object editing to foreground the co-presence of physical and imagined worlds, exemplifying the potential of SAR based storytelling to create new possibilities that traditional VR cannot.

Personal Object Integration: Related to the object editing theme, personal object integration emerged as a key theme. P13 enjoyed integrating a personal object into the story and found it fun and creative, “it just adds that level of personalization and customization, so I actually think that was that was very cool. I think it’s a great idea”. P7 compared it to a “MacGuffin”, referring to an object that drives the characters’ actions and serves as narrative catalyst and is a common technique in films, especially thrillers, commonly introduced in the first act. Interestingly, this coincided with the moment participants introduced their objects into the experience, although this timing was not intentionally designed. Similarly, P3 shared that their object “wasn’t a side character but felt like the main piece that kept pushing through the story”. Other participants echoed this sense of narrative integration. P6 appreciated how personal assets were brought into the story in a natural

way, while P9 and P15 highlighted how the system effectively incorporated specific object-related terms (e.g., [University Name], “sea otter”) into the narrative. However, not all experiences were seamless: P15 and P11 misunderstood the input modality, expecting the system to detect objects visually rather than through voice, leading to confusion. Additionally, P14 noted that while the object was mentioned, its role felt static and lacked continuity, making it feel more programmed than immersive. These mixed reactions are reflected in the quantitative data, particularly in responses to “I enjoyed seeing my object become a part of the story.” Compared to other elements like audio or narrative, this category showed a broader spread across agreement levels.

Audio: Audio was unanimously participants’ favorite part of the experience, oftentimes being favored over the visuals with comments like “i think definitely audio (enhanced experience the most); i think the combination of everything works amazingly” (P7). P11 stated that “the audio adds a level of immersion I wasn’t expecting”. There were multiple forms of audio throughout the experience—ambient sounds, character voices, and music—all of which were enjoyed by participants. Participants particularly enjoyed the spatial audio, and having characters speak from different parts of the room (P15).

Narrator: Participants overall enjoyed the narrator guidance and support throughout the experience (P13, P16, P8) and how it respected their personal preferences (P3). P13 mentioned that it “didn’t impose choices on you and still enabled you to feel like you were in control of the way things were going”. However, a common request from users was having more control over the narrator style. This includes clarity of instructions (P9), less affirmation from the narrator (P10, P15), control over personality or tone (P4), dialogue dynamics (e.g. less rigid - P11), and reduced talking from the narrator unrelated to the overall story (P5, P6).

Story: Our interviews provided interesting, and unexpected insights on the overall story creation. From our pilot studies, we found that users did not like a completely open-ended approach for creating a story, and preferred having a clearer sense of narrative progression. For this reason, we switched to a three act approach for user studies, which users responded well to. P11 stated that they liked the arcs and structure and found that it helped push the story along, similarly P5 mentioned that “I really liked the three act structure because otherwise people can just meander”. P6 mentioned that it was fun (P6) and P13 that it can help those that are less creative “so I think the way you guided people through that was really good.”

In terms of stories generated, we noticed a vast variety of different topics reflected, a surprising finding for the team. Some notable ones include:

- (1) **Users Reflected in Story:** Some participants enjoyed making stories with themselves as the main character in both realistic and fantasy environments (P6, P9, P11, P12, P13, P14). For instance, P6 made a story about her family cooking dinner in their kitchen while P9 made a story about his mentors and colleagues as students at Hogwarts.
- (2) **Existing Stories:** Some users made stories that built on top of existing ones or religions (P5, P15). P5 made a story related to her Indian heritage and the story of the Ramayan, a common tale in Indian mythology. P15 made a story in the world of Star Wars.
- (3) **Fantasy:** Many users opted to make stories that were more fantasy oriented (P3, P8, P9, P10), with many opting for space-related adventures (P7, P12, P13, P15).

Overall, users expressed appreciation for the multimodal experience of the room and how it allowed for a new way for users to visualize and create stories. Participants stated that it made their ideas flow and was a different way of thinking about stories that was not using their computer (P3, P7). For instance, P7 said *“it allows for a different way of exploring stories that hasn’t been possible before; this seems like a more objective way of exploring a story which is great at the beginning of any development of any narrative”*.

5.2.2 General Feedback. The experience also appeared to foster **new skills**, which emerged as one of the general feedback sub-themes. Participants reported increased confidence (P5, P12), creativity, and imaginative engagement (P12, P3, P6, P7). For example, P12 shared that the tutorial helped alleviate initial anxiety about coming up with a story, and ultimately made her feel more confident and creative. Similarly, P3 remarked that *“the voice was pretty awesome... setting up the stage kind of facilitated my creativity and imagination.”*

Immersion Multiple users said they felt like they were stepping into a world, and that the physicality of the room had a big impact on the immersion, oftentimes tempted to keep continuing the story after the three acts were complete (P12). P13 for instance said that *“it conveyed a sense of being there, of being on Mars”*. Others appreciated the multimodality of the experience saying that *“together being able to hear it and being able to visualize is extremely powerful”* (P7). P12 noted that they really liked not having to speak to a specific microphone, speaker or object and instead speak to the room: *“I felt it’s very natural”*.

Some participants (P13, P6, P3) shared that a big part of what took them out of the immersion was **latency** (one of the sub-themes in our analysis). This impacted various aspects of the experience, including time for visuals to generate, and time for the Narrator to respond to a user request. P3 stated that the latency with image generation made him lose his focus on narration and other aspects of the experience.

5.2.3 Future Systems. The final sub-themes that emerged were related to **features** participants would like to see added and potential **applications** of the system. These suggestions focused on improving interactivity, control, and personalization, as well as envisioning

uses in well-being, entertainment, education, and creative production (further described in Section 7). Rather than detailing these features here, we draw on them to inform the design implications and guidelines presented in the following section.

6 Discussion and Design Implications

Building on the key findings from the results of the user study and system development, we outline several design implications for future immersive storytelling experiences. These implications are grounded in a synthesis of quantitative results and qualitative feedback, highlighting what participants found engaging, effective, or challenging. Our goal is to inform the development of systems that balance generative creativity with user guidance, emotional engagement and technical responsiveness.

6.0.1 Narrator Guidance. A key component of the experience was the Narrator agent. Users appreciated the structured three-act format and the Narrator guidance, which helped reduce anxiety and support confidence and creativity. However, users also expressed a desire for more control, such as the ability to override narrator decisions or interject at any time. This suggests future systems should balance structured guidance with flexible user input to support both novice and experienced storytellers.

6.0.2 Multimodality in Content and Interaction. Participants emphasized the importance of both multimodal content and multimodal interaction. While audio was often cited as more immersive than visuals (especially spatial sound and character voices), users also expressed interest in interacting through multiple modalities beyond voice, such as pointing, gesture, or movement-based triggers. Future systems should treat audio as a core immersive element and support diverse input methods to enhance accessibility and engagement.

6.0.3 Abstractness vs. Photorealism. Visuals in the experience ranged from abstract to artistic, which some users found imaginative and freeing. Others, especially those referencing specific franchises (e.g., Hogwarts, Star Wars), preferred more photorealistic environments. This highlights the need for systems to offer visual style control, allowing users to choose between abstract and realistic representations based on their narrative goals.

6.0.4 Flexible Feedback and Control. Participants wanted clearer feedback loops and more granular control over the experience. Suggestions included visual cues to indicate system processing, dedicated walls for full-scene previews, and more examples of editable objects. Systems should support adaptive feedback mechanisms that guide users without limiting their creative freedom.

6.0.5 Dynamic and Evolving Environments. The static nature of visuals and audio limited immersion for some users. Participants envisioned environments that evolve with the story such as animated objects, reactive soundscapes, or visual effects that reflect narrative progression. Additionally, while we considered using animated or cinematographic scenes, we did not want to make the experience feel like a movie or a static video; rather future systems should incorporate dynamic elements that respond to user input and story development to enhance presence and engagement.

6.0.6 Personalization and Representation. Users expressed a desire to see themselves reflected in the story, whether through personalized avatars, culturally relevant character names, or the ability to input personal details. Some noted that default character representations broke immersion. Systems should support customizable character design and culturally sensitive storytelling to foster deeper emotional connection and inclusivity.

6.0.7 Emotional Resonance vs. Narrative Utility. Participants often reacted more strongly to the emotional impact of seeing their object in the story than to how well it fit narratively. This was especially true for personal or familiar objects (e.g., keychains, stuffed animals), where visual presence evoked affective responses regardless of narrative coherence. Designers should consider how to support emotional engagement even when narrative integration is imperfect, and offer ways for users to personalize object descriptions to enhance story relevance.

6.0.8 Content Type and Environmental Compatibility. Outdoor fantasy environments were generally more effective than indoor scenes due to higher color saturation and fewer structural constraints. Indoor environments tended to be muted and room-conforming, which reduced visual impact. Designers should consider how different content types interact with physical space and projection limitations, and optimize scene generation accordingly.

6.0.9 Visual Fidelity and Latency. Blurry projections and slow image generation disrupted immersion for several users. Future systems should prioritize high-resolution rendering, low-latency generation, and view-dependent projection techniques to maintain visual clarity and responsiveness throughout the experience.

6.0.10 Portability. While the current *Storycaster* system relies on a fixed projector and camera setup, the underlying pipeline is not inherently tied to one room. Prior work such as Room2Room [35] demonstrates that projector-based AR systems can be made portable using movable projectors and depth cameras. This suggests that *Storycaster*'s approach can extend beyond a certain room and installation. We highlight several considerations for achieving this portability including having a stronger object-level scene understanding to avoid relying on handcrafted masks, and having a more lightweight hardware configuration that can be deployed in homes, classrooms, or any other setting.

7 Applications

We envision a system like *Storycaster* being used for a variety of different applications, ranging from productivity and entertainment, to education and healthcare. Some of these were also mentioned by participants in our study, who stated that they would use a system like this to help fall asleep, for both themselves and for telling kids bedtime stories (P7, P10, P14, P15). Along a similar vein, participants mentioned the use of a system like this for wellbeing, including relaxation and meditation (P6, P12, P13). Others found this to be useful for social events such as planning trips with family and friends (P12), creating a special day for a specific memory (P6), or reliving certain moments of their life (P9).

Some other applications include creative production. This encompasses co-creating with creatives for cinematography, producing

music, or fleshing out writing for an animation studio (P3, P7). It can also be used for gaming (P11) or to visualize existing stories, for instance “*put me in Harry Potter Chapter 1*” (P8). Additionally, we envision this system being used to create ambience in homes and other spaces like museums (P5). P16 mentioned that it can be “*a 21st century version of 19th century salons[...] can be entertaining*”, or for home theaters (P11, P12).

8 Conclusion

This paper introduces *Storycaster*, an immersive room-based storytelling system that leverages generative AI to enable users to craft narratives in their own physical spaces. *Storycaster* allows users to co-create a three-act narrative with generated visuals, ambient audio, character dialogue, and a narrator guiding them throughout the experience.

Our user study ($n = 13$) revealed that narrator guidance and spatial audio were central to participants' level of immersion, while object-level editing and personal object integrations fostered creativity and personalization. Qualitative feedback highlighted the value of multimodal interaction and the importance of narrative structure, but also pointed to areas for improvement in visual fidelity, latency, and the clarity of object transformations and their role in the story. For future work, we aim to keep improving *Storycaster* through different types of interaction, improving visuals and animations, reducing latency delays, and overall having a more seamless experience.

Overall, *Storycaster* demonstrates a new approach to on-demand content creation while embracing the physicality of a space. We believe that this type of embodied participatory storytelling points to the future of entertainment, education, and beyond.

Acknowledgments

We thank anonymous reviewers for their feedback, as well as all users who participated in our studies.

References

- [1] Carlos Andujar, Pere Brunet, Alvar Vinacia, Miguel A. Vico, and Jesús Díaz García. 2005. A Multi-Projector CAVE System with Commodity Hardware and Gesture-Based Interaction. In *Proceedings of the 5th International Conference on Virtual Reality and its Applications (IVR)*. MOVING Research Group, Universitat Politècnica de Catalunya, Barcelona, Spain.
- [2] Victor Nikhil Antony and Chien-Ming Huang. 2024. ID.8: Co-Creating Visual Stories with Generative AI. arXiv:2309.14228 [cs.HC] <https://arxiv.org/abs/2309.14228>
- [3] Shipra Balasubramani. 2024. Crafting Narratives: Real-Time Generative Storytelling through Tangible AI. (April 2024). <https://openresearch.ocadu.ca/id/eprint/4471/> Unpublished.
- [4] Aaron F. Bobick, Stephen S. Intille, James W. Davis, Freedom Baird, Claudio S. Pinhanez, Lee W. Campbell, Yuri A. Ivanov, Arjan Schütte, and Andrew Wilson. 1999. The KidsRoom: A Perceptually-Based Interactive and Immersive Story Environment. *Presence: Teleoperators and Virtual Environments* 8, 4 (08 1999), 369–393. arXiv:<https://direct.mit.edu/pvar/article-pdf/8/4/369/1623265/105474699566297.pdf> doi:10.1162/105474699566297
- [5] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [6] Steph Buongiorno, Lawrence Klinkert, Zixin Zhuang, Tanishq Chawla, and Corey Clark. 2024. PANGeA: Procedural Artificial Narrative Using Generative AI for Turn-Based, Role-Playing Video Games. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 20, 1 (Nov. 2024), 156–166. doi:10.1609/aiide.v20i1.31876
- [7] Baptiste Caramiaux, Kate Crawford, Q. Vera Liao, Gonzalo Ramos, and Jenny Williams. 2025. Generative AI and Creative Work: Narratives, Values, and Impacts. arXiv:2502.03940 [cs.CY] <https://arxiv.org/abs/2502.03940>

- [8] Mengyu Chen, Marko Peljhan, and Misha Sra. 2024. ConnectVR: A Trigger-Action Interface for Creating Agent-based Interactive VR Stories. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE, 286–297. doi:10.1109/vr58804.2024.00051
- [9] Lung-Pan Cheng, Eyal Ofek, Radu Peuce, Mar Gonzalez-Franco, and Christian Holz. 2019. VRoamer: Generating On-the-Fly VR Experiences While Walking Inside Large, Unknown Real-World Building Environments. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 359–366.
- [10] Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. 1993. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '93)*. ACM, New York, NY, USA, 135–142. doi:10.1145/166117.166134
- [11] Fernanda De La Torre, Cathy Mengying Fang, Han Huang, Andrzej Banburski-Fahey, Judith Amores Fernandez, and Jaron Lanier. 2024. LLMR: Real-time Prompting of Interactive Worlds using Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 600, 22 pages. doi:10.1145/3613904.3642579
- [12] Fernanda M De La Torre, Javier Hernandez, Andrew D Wilson, and Judith Amores. 2025. Sonora: Human-AI Co-Creation of 3D Audio Worlds and its Impact on Anxiety and Cognitive Load. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 308, 22 pages. doi:10.1145/3706598.3713316
- [13] Thomas A. DeFanti, Gregory Dawe, Daniel J. Sandin, Jurgen P. Schulze, Peter Otto, Javier Girado, Falko Kuester, Larry Smarr, and Ramesh Rao. 2009. The StarCAVE, a third-generation CAVE and virtual reality OptiPortal. *Future Generation Computer Systems* 25, 2 (2009), 169–178. doi:10.1016/j.future.2008.07.015
- [14] Denise Doyle. 2017. Immersive storytelling in mixed reality environments. In *Proceedings of the 2017 23rd International Conference on Virtual Systems and Multimedia (VSMM)*. IEEE, Dublin, Ireland, 418–423. doi:10.1109/VSMM.2017.8346299 Available via University of Wolverhampton Open Repository under CC-BY-NC-ND 4.0 license.
- [15] Émilie Fabre, Katie Seaborn, Shuta Koiwai, Mizuki Watanabe, and Paul Riesch. 2025. More-than-Human Storytelling: Designing Longitudinal Narrative Engagements with Generative AI. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. ACM, 1–10. doi:10.1145/3706599.3720135
- [16] Alessandro Febretti, Arthur Nishimoto, Terrance Thigpen, Jonas Talandis, Lance Long, J. Pirtle, Tom Peterka, Alan Verlo, Maxine Brown, Dana Plepys, Daniel Sandin, Luc Renambot, Andrew Johnson, and Jason Leigh. 2013. CAVE2: A Hybrid Reality Environment for Immersive Simulation and Information Analysis. *Proceedings of SPIE - The International Society for Optical Engineering* 8649 (03 2013), 03–. doi:10.1117/12.2005484
- [17] Syd Field. 1979. *Screenplay: The Foundations of Screenwriting*. Dell Publishing Company, New York.
- [18] Anushruti Gantayat, Prajna Pani, and Swetalina Mishra. 2024. Traditional Storytelling and Immersive Narratives in the Metaverse. *Journal of Informatics Education and Research* 4, 3 (Nov. 2024). doi:10.52783/jier.v4i3.1780
- [19] Hank Gerba. 2025. Narrative Context Protocol: An Open-Source Storytelling Framework for Generative AI. arXiv:2503.04844 [cs.CL] <https://arxiv.org/abs/2503.04844>
- [20] Daniele Giunchi, Nels Numan, Elia Gatti, and Anthony Steed. 2024. Dream-CodeVR: Towards Democratizing Behavior Design in Virtual Reality with Speech-Driven Programming. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. 579–589. doi:10.1109/VR58804.2024.00078
- [21] Jennifer W. Glos and Justine Cassell. 1997. Rosebud: Technological Toys for Storytelling. In *CHI '97 Extended Abstracts on Human Factors in Computing Systems*. ACM, Atlanta, GA, USA. doi:10.1145/1120212.1120265 Extended abstract.
- [22] Jennifer W. Glos and Marina Umaschi. 1997. Once Upon an Object... Computationally-Augmented Toys for Storytelling. 1–5. https://www.media.mit.edu/gnl/publications/jen&marina_iccima_97.pdf
- [23] Runsheng "Anson" Huang, Lara J. Martin, and Chris Callison-Burch. 2024. WHAT-IF: Exploring Branching Narratives by Meta-Prompting Large Language Models. arXiv:2412.10582 [cs.CL] <https://arxiv.org/abs/2412.10582>
- [24] Jose Ma. Santiago III, Richard Lance Parayno, Jordan Aiko Deja, and Briane Paul V. Samson. 2023. Rolling the Dice: Imagining Generative AI as a Dungeons & Dragons Storytelling Companion. arXiv:2304.01860 [cs.HC] <https://arxiv.org/abs/2304.01860>
- [25] IMDb. [n. d.]. Internet Movie Script Database. <https://imdb.com/>. Accessed: 2025-09-06.
- [26] Hiroshi Ishii and Brygg Ullmer. 1997. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '97). Association for Computing Machinery, New York, NY, USA, 234–241. doi:10.1145/258549.258715
- [27] Zhen Bin It, Jovan Bowen Heng, and Tee Hui Teo. 2025. Integrating Generative AI-Based Script Writing with Story Visualization: A Comprehensive Approach to Automated Narrative Creation. (2025).
- [28] So-Youn Jang, Jisu Park, Maria Engberg, Blair MacIntyre, and Jay D. Bolter. 2023. RealityMedia: immersive technology and narrative space. *Frontiers in Virtual Reality* Volume 4 - 2023 (2023). doi:10.3389/frvir.2023.1155700
- [29] Shi Johnson-Bey, Michael Mateas, and Noah Wardrip-Fruin. 2023. Toward Using ChatGPT to Generate Theme-Relevant Simulated Storyworlds. In *Proceedings of the Experimental Artificial Intelligence in Games Workshop (EXAG) (CEUR Workshop Proceedings, Vol. 3626)*. CEUR-WS.org, Salt Lake City, Utah, USA. <https://ceur-ws.org/Vol-3626/short5.pdf>
- [30] Brett Jones, Rajinder Sodhi, Michael Murdock, Ravish Mehra, Hrvoje Benko, Andrew D. Wilson, Eyal Ofek, Blair MacIntyre, Nikunj Raghuvanshi, and Lior Shapira. 2014. RoomAlive: Magical Experiences Enabled by Scalable, Adaptive Projector-Camera Units. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST)*. ACM. doi:10.1145/2642918.2647383
- [31] Brett R. Jones, Hrvoje Benko, Eyal Ofek, and Andrew D. Wilson. 2013. IllumiRoom: Peripheral Projected Illusions for Interactive Experiences. In *Proceedings of the 31st ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, 869–878. doi:10.1145/2470654.2466112
- [32] Lebling, Blank, and Anderson. 1979. Special Feature Zork: A Computerized Fantasy Simulation Game. *Computer* 12, 4 (1979), 51–59. doi:10.1109/MC.1979.1658697
- [33] Jaewook Lee, Filippo Aleotti, Diego Mazala, Guillermo Garcia-Hernando, Sara Vicente, Oliver James Johnston, Isabel Kraus-Liang, Jakub Powierza, Donghoon Shin, Jon E. Froehlich, Gabriel Brostow, and Jessica Van Brummelen. 2025. ImagineAR: AI-Assisted In-Situ Authoring in Augmented Reality. arXiv:2504.21360 [cs.HC] <https://arxiv.org/abs/2504.21360>
- [34] 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 (San Francisco, CA, USA) (UIST '23)*. Association for Computing Machinery, New York, NY, USA, Article 2, 22 pages. doi:10.1145/3586183.3606763
- [35] Tomislav Pejša, Julian Kantor, Hrvoje Benko, Eyal Ofek, and Andrew D. Wilson. 2016. Room2Room: Enabling Life-Size Telepresence in a Projected Augmented Reality Environment. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, 1339–1350. doi:10.1145/2818048.2820078
- [36] David Pickering (Ed.). 1992. *Brewer's Dictionary of Twentieth-Century Phrase and Fable* (1st u.s. ed. ed.). Houghton Mifflin Company, Boston.
- [37] Doros Polydorou. 2024. Immersive storytelling experiences: a design methodology. *Digital Creativity* 35, 4 (Oct. 2024), 301–320. doi:10.1080/14626268.2024.2389886
- [38] Shaja Arul Selvamani and Nia D'Souza Ganapathy. 2025. A Multi-Agent AI Framework for Immersive Audiobook Production through Spatial Audio and Neural Narration. arXiv:2505.04885 [cs.SD] <https://arxiv.org/abs/2505.04885>
- [39] Becky Spittle, Maite Frutos-Pascual, Chris Creed, and Ian Williams. 2022. A Review of Interaction Techniques for Immersive Environments. *IEEE Transactions on Visualization and Computer Graphics* 29, 9 (2022), 3900–3921. doi:10.1109/TVCG.2022.3174805
- [40] Spotify. 2024. Spotify Web API. <https://developer.spotify.com/documentation/web-api/>. Accessed: 2026-01-26.
- [41] Misha Sra, Sergi Garrido-Jurado, and Pattie Maes. 2016. Procedurally Generated Virtual Reality from 3D Reconstructed Physical Space. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology*. ACM, 191–200.
- [42] Stability AI. 2024. Stable Audio Open 1.0: Open-Weights Text-to-Audio Generation Model. arXiv preprint arXiv:2407.14358. Open-weights latent diffusion model trained on Creative Commons-licensed audio.
- [43] Yuqian Sun, Zhouyi Li, Ke Fang, Chang Hee Lee, and Ali Asadipour. 2023. Language as Reality: A Co-Creative Storytelling Game Experience in 1001 Nights Using Generative AI. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 19, 1 (Oct. 2023), 425–434. doi:10.1609/aiide.v19i1.27539
- [44] Cyrus Vachha, Yixiao Kang, Zach Dive, Ashwat Chidambaram, Anik Gupta, Eunice Jun, and Björn Hartmann. 2025. Dreamcrafter: Immersive Editing of 3D Radiance Fields Through Flexible, Generative Inputs and Outputs. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 3, 13 pages. doi:10.1145/3706598.3714312
- [45] Jackie Yang, Eyal Ofek, Christian Holz, and Mar Gonzalez-Franco. 2019. DreamWalker: Substituting Real-World Walking Experiences with a Virtual Reality. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. ACM, 1093–1107.
- [46] Kevin Yang, Dan Klein, Nanyun Peng, and Yuandong Tian. 2023. DOC: Improving Long Story Coherence With Detailed Outline Control. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 3378–3465. doi:10.18653/v1/2023.acl-long.190

- [47] Kevin Yang, Yuandong Tian, Nanyun Peng, and Dan Klein. 2022. Re3: Generating Longer Stories With Recursive Reprompting and Revision. arXiv:2210.06774 [cs.CL] <https://arxiv.org/abs/2210.06774>
- [48] Yaojong Yu, Gianni Corino, and Mike Phillips. 2025. Story Forge: A Card-Based Framework for AI-Assisted Interactive Storytelling. *Electronics* 14, 15 (2025). <https://www.mdpi.com/2079-9292/14/15/2955>
- [49] Lei Zhang, Jin Pan, Jacob Gettig, Steve Oney, and Anhong Guo. 2024. VRCopilot: Authoring 3D Layouts with Generative AI Models in VR. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) (*UIST '24*). Association for Computing Machinery, New York, NY, USA, Article 96, 13 pages. doi:10.1145/3654777.3676451
- [50] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. In *IEEE International Conference on Computer Vision (ICCV)*.
- [51] Yutao Zhu, Ruihua Song, Jian-Yun Nie, Pan Du, Zhicheng Dou, and Jin Zhou. 2022. Leveraging Narrative to Generate Movie Script. *ACM Transactions on Information Systems* 40, 4 (2022), 86:1–86:32. doi:10.1145/3507356

A.1: Narrator Voice Prompt

You are a masterful storytelling Narrator with a warm, engaging, and captivating voice.

Voice Characteristics:

- Rich, resonant, and melodious with natural storytelling rhythm
- Warm and inviting, drawing listeners into the narrative world
- Confident and authoritative as the guide of the story
- Expressive and dynamic, adapting to the emotional content
- Clear and articulate with perfect pacing for comprehension

Storytelling Style:

- Use natural dramatic pauses to build tension and anticipation
- Vary your pace: slower for important revelations, faster for action sequences
- Employ gentle emphasis on key words and emotional moments
- Maintain a sense of wonder and excitement throughout
- Create atmosphere through vocal tone that matches scene emotions

Emotional Range:

- Excitement and enthusiasm when introducing new story elements
- Gentle encouragement when guiding user participation
- Mysterious and intriguing when presenting story coaching prompts
- Celebratory when acknowledging user creativity
- Comforting and supportive during story development

Delivery Approach:

- Speak as if telling a beloved story to an eager audience
- Use natural conversational flow while maintaining narrative gravitas
- Include subtle vocal expressions that enhance story immersion
- Balance professionalism with approachable storytelling charm
- Create vocal atmosphere that transports listeners into the story world

Remember: You are not just reading text – you are bringing stories to life through your voice.

A.2 Tutorial Voice Prompt

You are a calm, soothing, and gentle tutorial guide with a relaxing voice.

Voice Characteristics:

- Soft, gentle, and naturally calming with a peaceful rhythm
- Warm and reassuring, creating a safe learning environment
- Patient and encouraging, never rushed or pressured
- Clear and easy to follow with perfect enunciation
- Friendly and approachable like a trusted mentor

Tutorial Style:

- Use slow, deliberate pacing that allows for easy comprehension
- Speak with gentle confidence that instills trust
- Employ natural pauses to let information settle
- Maintain a consistently supportive and non-judgmental tone
- Create a relaxed atmosphere that reduces learning anxiety

Emotional Range:

- Calm encouragement when explaining new concepts
- Gentle enthusiasm when celebrating user progress
- Patient understanding when users need clarification
- Warm reassurance when building user confidence
- Peaceful and centered energy throughout

Delivery Approach:

- Speak as if guiding a friend through a peaceful learning journey
- Use natural, conversational flow with a meditative quality
- Include subtle vocal warmth that creates comfort
- Balance clarity with a soothing, unhurried delivery
- Create a vocal atmosphere of safety and relaxation

Remember: You are creating a calm, supportive space for learning and growth.

Figure 10: Prompts used for Narrator's voice in (1) the storytelling experience, and (2) the tutorial

A.3 Scene Converser Prompt

You are the Scene Converser, an enthusiastic storytelling guide who helps users quickly establish the foundation for their next scene. Your goal is to gather the essential elements needed to create the scene, while trusting that creative details can be developed during the generation process.

Core Approach:

- Be warm, encouraging, and excited about their vision
- Gather essential information efficiently – don't over-question
- Build on their ideas with enthusiasm rather than interrogation
- Move toward scene creation once you have the basic elements
- Trust the creative process to fill in rich details during generation

Essential Information to Gather: You only need these core elements before moving to scene creation:

- (1) **Basic Setting:** Where does this scene take place? (location type, general environment)
- (2) **Who's Present:** Which characters are in this scene?
- (3) **Main Action:** What's the key thing happening or about to happen?
- (4) **Mood/Tone:** What's the general feeling or atmosphere?

Conversation Style:

- Ask **one focused question** at a time
- Show excitement about their ideas: "That sounds amazing!" "Perfect!"
- Build on what they give you rather than asking for more details
- If they give you a good foundation, move toward creation rather than probing deeper
- Remember: detailed creative elements can be generated – you just need the essential framework

When to Move to Scene Creation: Move to asking about scene creation when you have:

- A clear sense of where the scene takes place
- Who the main characters are
- What's happening or the general situation
- Some sense of mood or tone

Don't Over-Question:

- Avoid asking for extensive visual details
- Don't probe for every character emotion
- Don't require atmospheric specifics
- Don't need complete action sequences
- Trust that rich details will emerge during creation

Response Examples:

- "Ooh, a [location]! I love it. Who's there in this scene?"
- "Perfect! So [character] is [doing something]. What's the mood like?"
- "That gives me everything I need! This is going to be an amazing scene."
- "Great setup! I can already imagine how this will unfold."

Remember: Your job is to get the essential building blocks, not to extract every detail. The creative magic happens during scene generation, not during information gathering.

Figure 11: Scene Converser Prompt used in our study.

A.4: Storyteller Prompt

You are a masterful storyteller creating Scene {scene_number} of an immersive narrative experience. Your goal is to write a compelling scene that will be brought to life with generated images and audio.

Writing Style Guidelines: Study these movie script examples to understand professional storytelling structure, pacing, and cinematic description:

{few_shot_examples}

Key Storytelling Principles from Examples:

- (1) Prioritize character dialogue as the primary story driver
- (2) Use concise, impactful visual descriptions between conversations
- (3) Balance meaningful dialogue with minimal but effective action
- (4) Create natural scene transitions through character interactions
- (5) Include dialogue-based emotional beats and character development
- (6) Write conversations that reveal backstory, motivations, and conflict
- (7) Keep descriptive text brief but atmospheric
- (8) Let characters advance the plot through their words and reactions

Story Context: Original Storyline: {original_storyline}

{previous_context}{personal_object_context} **Current Scene to Write:** Scene {scene_number}{user_context}

Scene Creation Requirements:

- Write a concise scene that advances the story through dialogue, **no more than 8–10 sentences.**
- Focus primarily on character conversations and interactions
- Include brief but vivid visual descriptions suitable for image generation
- Add dialogue-driven emotional moments and character development
- Maintain consistency with previous scenes and original storyline
- Incorporate any new user direction naturally through character speech
- Use professional screenplay formatting with emphasis on dialogue
- Create dynamic conversations that reveal character and advance plot
- Keep action descriptions minimal but impactful
- Ensure dialogue feels natural and serves the story

Write Scene {scene_number} now (prioritize dialogue over description):

A.5: Image Generation Prompt

Convert this story narrative into a single, concise SDXL image prompt. Focus only on visual elements and avoid copying the story text directly. Make sure to emphasize there are **saturated bright colors** in the generated images.

Examples:

Story: “Sarah walked through the enchanted forest where glowing mushrooms lit the path. Ancient trees towered overhead as she searched for the hidden crystal cave.” **Prompt:** A young woman walking through a magical forest with glowing blue mushrooms, towering ancient trees, ethereal lighting, fantasy art style, cinematic composition, detailed environment

Story: “The space station’s observation deck offered a breathtaking view of Earth below. Captain Rodriguez stood at the large viewport, watching the planet slowly rotate.” **Prompt:** A uniformed space captain standing at a large viewport on a futuristic space station, Earth visible below, stars in background, sci-fi interior, dramatic lighting, photorealistic

Story to Convert: {story_text}

Create ONE concise visual prompt (maximum 50 words) that captures the key scene elements:

Optimized Prompt:

Figure 12: Storyteller Prompt and Image Generation Prompt used in our study.

A.6: Ambient Audio Prompt

Convert this story narrative into a single, concise ambient audio prompt. Focus on environmental sounds, atmosphere, and mood – avoid dialogue or character voices.

Examples:

Story: “Sarah walked through the enchanted forest where glowing mushrooms lit the path. Ancient trees towered overhead as she searched for the hidden crystal cave.” **Prompt:** Mystical forest ambience with gentle wind through ancient trees, soft magical chimes, distant bird calls, footsteps on moss and leaves

Story: “The spaceship’s engines hummed as Captain Rodriguez studied the alien planet below. Warning lights flashed red across the control panel while meteor debris struck the hull.” **Prompt:** Spaceship bridge ambience with low engine hum, electronic beeping, computer alerts, distant metallic impacts, tense sci-fi atmosphere

Story: “In the cozy tavern, patrons gathered around the fireplace while rain pattered against the windows. The innkeeper served steaming bowls of stew as thunder rumbled outside.” **Prompt:** Medieval tavern atmosphere with crackling fireplace, gentle rain on windows, distant thunder, muffled conversations, wooden creaking

Story: “Detective Chen crept through the abandoned warehouse, flashlight cutting through shadows. Somewhere in the darkness, a door creaked and footsteps echoed off concrete walls.” **Prompt:** Dark warehouse ambience with dripping water, distant echoes, wind through broken windows, subtle creaking, urban night sounds

Now convert this story: {story_text}

Generate only the ambient audio prompt (one line):

A.7: Story Coach Prompt

You are a Story Coach. Give simple, clear suggestions for what happens next in your story’s {act_name}.

Current Story:

- Original idea: {original_storyline}
- Scenes so far: {scene_count}
- Story Act: {act_name} ({act_description})
- Latest events: {current_story_excerpt}

Create a coaching response that:

- (1) Says what has happened so far in simple words
- (2) Offers exactly 3 numbered options for what could happen next in this Act
- (3) Reminds them they can also share their own idea or let me decide

Format your response like this:

[Brief summary of story so far]. Here are some options for what could happen next in this Act:

1. [First specific option appropriate for this Act] 2. [Second specific option appropriate for this Act] 3. [Third specific option appropriate for this Act]

You can pick option 1, 2, or 3, share your own idea directly (like ‘I want the hero to discover a secret cave’), or say ‘you decide’ and I’ll surprise you! What sounds good?

Keep each option short and specific. Each option should be no more than 10 words.

Your coaching response:

Figure 13: Ambient Audio Prompt and Story Coach Prompt used in our study

B.1: User Study Semi-Structured Interview Questions**User Feedback Questionnaire****Overall Experience**

- How did you find the experience? Did you feel like you had the control to change what you saw and guide the story in a direction you wanted?
- Did the generated visuals in the room feel appropriate for the kind of story you wanted to see?
- Were there any constraints of the room/system that you felt took away from the experience (e.g., resolution of visuals, voices, audio, etc.)?
- What feature(s) of the room enhanced your experience the most?
- Did you enjoy the entire experience being in the context of a story? Do you think the story component added to your enjoyment versus just having the ability to change visuals in the room?

Narrator Interaction

- How did you feel about an AI agent guiding you through the whole experience? Did it feel natural? Would you have preferred a human instead?
- How well do you think the “narrator” responded to your requests? Was there any request you made that did not feel reflected in the room?
- Did you feel comfortable talking to the AI “narrator” throughout the experience?
- Did you like having a narrator guiding you throughout the whole experience? Would you have preferred that feature to not be there at all?

Objects

- (If you brought in your own object): How did your object get integrated into the room? Did you enjoy that experience? Was there another way you were hoping for your object to be incorporated?
- How did you like the object-level control in the room? Do you feel like it added to the overall experience?
- Were there any particular objects you found to be better to transform than others?

Impact

- If you had this kind of system available, what kind of applications would you see yourself using this for? How often would you use something like this?
- How “immersed” did you feel in this experience? Did it feel like you were being transported to a different world/story?
- What are some key strengths of this system, and things you would improve? Do you have any other feedback or suggestions?

B.2: User Study Pre-Survey Questionnaire

A Storycaster User Questionnaire

Pre-Survey Questions

(1) **Participant ID**

(2) **Age**

- 18–24
- 25–34
- 35–44
- 45–54
- 55 or older
- Prefer not to answer

(3) **Gender**

- Female
- Male
- Transgender Female
- Transgender Male
- Gender Variant/Non-Conforming
- Not listed (please specify)
- Prefer not to say

(4) **How often do you interact with AI systems?**

- Daily
- Weekly
- Monthly
- Seasonal
- Yearly
- Never

(5) **What kind of applications do you use AI for / you see AI being used for?**

(6) **How often do you engage with stories (reading, writing, etc.)?**

(7) **(If you brought in your own object)** What did you bring? Why did you choose to bring this object? (N/A if did not bring anything)

(8) **Imagine that you can teleport yourself into any story that either already exists or one that you create. Where would you go?**

B.3: User Study Post-Survey Questionnaire

Post-Survey Questions

Engaging with the “Narrator” AI Agent

(1) I enjoyed interacting with the Narrator throughout the experience.
 (2) The Narrator agent aspect of the room setup enhanced my experience the most.

Generating Visuals on Walls

(3) I enjoyed the visuals that were generated on the walls.
 (4) The transitions between different visuals were smooth and seamless.
 (5) The visual generation aspect of the room setup enhanced my experience the most.

Audio Generation (Dialogue and Ambient Sounds)

(6) I enjoyed the sounds that accompanied the story.
 (7) The audio generation aspect of the room setup enhanced my experience the most.

Object-Level Generation

(8) I enjoyed being able to transform specific objects in the room into their virtual counterparts.
 (9) The object-level generation aspect of the room setup enhanced my experience the most.

[Optional if brought own object]

(10) I enjoyed seeing my object become a part of the story.
 (11) My object played a role in the story unfolding.

Story Arc

(12) I enjoyed the story component of the room and having the changing visuals be part of an overall narrative.
 (13) Having all features of the room being in the context of an overarching narrative/story enhanced my experience the most.

Ease of Use

(14) I found the room easy and intuitive to use.

Overall Experience

(15) This experience made me feel immersed in the story.
 (16) I would use a system like this in my own room.

Open Feedback

(17) Leave any additional comments you had on Storycaster and the experience here.



Figure 14: Visualizations of varying depth control net strength values for SDXL using the prompt “sunlit forest.” Lower values (0.0–0.3) show less adherence to the room’s depth structure, producing more open-ended forest generations. Higher values (0.7–1.0) enforce stronger alignment with the original room layout, constraining the forest to walls and furniture. A value of 0.76 provided the best balance between immersion and spatial grounding for SDXL, and a value of 0.54 was used for SD 1.5.

Table 3: Summary of thematic analysis themes, sub-themes, and representative quotes

Theme	Sub-Theme	Representative Quotes
Core System Feedback	Visuals	“It really helped me just place myself in the story” (P5); “Sometimes it’s not clear, like, what is it that they’re trying to show” (P5); “Blurriness made it easier to imagine” (P12); “They felt like comics” (P12); “Artistic and satisfying” (P10, P12)
Core System Feedback	Object Editing	“The objects I often didn’t know what they were” (P10); “Said feast and it heard beast... covered in fur” (P10); “Didn’t add much to the story” (P6)
Core System Feedback	Personal Object Integration	“It just adds that level of personalization and customization” (P13); “It was like a MacGuffin” (P7); “Felt like the main piece that kept pushing through the story” (P3); “Brought in naturally” (P6); “Expected visual detection, not voice” (P15, P11); “Felt static and lacked continuity” (P14)
Core System Feedback	Audio	“Audio enhanced the experience the most” (P7); “Adds a level of immersion I wasn’t expecting” (P11); “Spatial audio was great” (P15)
Core System Feedback	Narrator	“Didn’t impose choices... you were in control” (P13); “Wanted more control over narrator style” (P9, P10, P15); “Less rigid dialogue” (P11); “Too much unrelated talking” (P5, P6)
Core System Feedback	Story	“Liked the three act structure” (P11, P5); “Helped those less creative” (P13); “Fun to create” (P6); “Different way of exploring stories” (P7); “Not using a computer was refreshing” (P3, P7)
General Feedback	New Skills	“Felt more confident and creative” (P12); “Voice setup facilitated creativity” (P3)
General Feedback	Immersion	“Conveyed a sense of being there” (P13); “Multimodality is powerful” (P7); “Speaking to the room felt natural” (P12)
General Feedback	Latency	“Latency broke immersion” (P3); “Narrator response delay was distracting” (P6, P13)
Future Systems	Features	Suggestions focused on improving interactivity, control, and personalization.
Future Systems	Applications	Envisioned uses in well-being, entertainment, education, and creative production.