

DiaryPlay: AI-Assisted Creation of Interactive Story Vignettes for Everyday Storytelling

Jiangnan Xu
Tampere University
Tampere, Finland
Rochester Institute of Technology
Rochester, New York, USA
jiangnan.xu@tuni.fi

Haeseul Cha
KAIST
Daejeon, Republic of Korea
jjchs1@kaist.ac.kr

Gosu Choi
KAIST
Daejeon, Republic of Korea
coregosu12@kaist.ac.kr

Gyu-cheol Lee
Korea Telecom
Seoul, Republic of Korea
gc.lee@kt.com

Yeo-Jin Yoon
Korea Telecom
Seoul, Republic of Korea
yjin.yun@kt.com

Zucheul Lee
Korea Telecom
Seoul, Republic of Korea
polelee@kt.com

Konstantinos Papangelis
Rochester Institute of Technology
Rochester, New York, USA
kxpigm@r.rit.edu

Dae Hyun Kim*
Yonsei University
Seoul, Republic of Korea
POSTECH
Pohang, Republic of Korea
dhkim16@yonsei.ac.kr

Juho Kim*
KAIST
Daejeon, Republic of Korea
SkillBench
Santa Barbara, CA, USA
juhokim@kaist.ac.kr

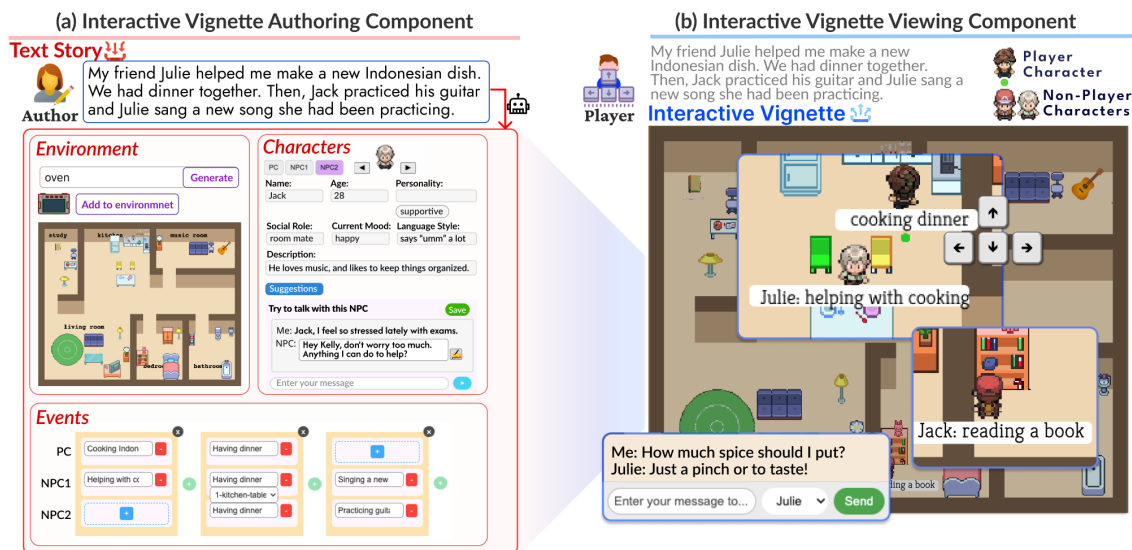


Figure 1: DIARYPLAY assists everyday storytellers in creating interactive vignettes. It contains two components: (a) The Interactive Vignette Authoring Component, which guides the everyday storyteller (author) to complete and refine the interactive vignette elements through surfacing the gap between what is captured by the natural language story input and the information requirements for an interactive vignette, (b) The Interactive Vignette Viewing Component, which enables the audience (player) to role-play the main character (player character; PC) to follow the author-defined storyline or engage in divergent activities within the environment, and plan non-player characters' (NPCs) behaviors in response to the player interactions.

* Co-responding authors



Abstract

An interactive vignette is a visual storytelling medium that lets the audience role-play a character and interact with non-player characters (NPCs) and the digital environment. Yet, the authoring complexity of interactive vignettes has obstructed their adoption in everyday storytelling, which builds on immediacy. We introduce DIARYPLAY, an AI-assisted authoring system that generates interactive vignettes from text stories. The *Authoring Component* visually elicits three core elements (environment, characters, events) through automation and author refinement. The *Viewing Component* delivers an interactive story to the audience using an LLM-powered *Controlled Divergence Module*, which allows divergent player and NPC behaviors within the boundaries defined by the author's intended story. A technical evaluation shows that the Controlled Divergence module generates believable NPC activities based on both character persona and storyline. A user study demonstrates that DIARYPLAY enables low-effort authoring of interactive vignettes for everyday storytelling while providing engaging viewing experiences and conveying the core story message.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI; Interactive systems and tools.**

Keywords

Interactive Storytelling, Everyday Storytelling, Interactive Vignette, Role-Playing, Authoring System, Human-AI Collaboration, Large Language Model

ACM Reference Format:

Jiangnan Xu, Haeseul Cha, Gosu Choi, Gyu-cheol Lee, Yeo-Jin Yoon, Zuchel Lee, Konstantinos Papangelis, Dae Hyun Kim, and Juho Kim. 2026. Diary-Play: AI-Assisted Creation of Interactive Story Vignettes for Everyday 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, 21 pages. <https://doi.org/10.1145/3772318.3790572>

1 Introduction

Everyday storytelling is an intrinsic part of human nature [35]. People not only tell stories about daily moments verbally to those around them, but also share everyday stories digitally on social media platforms (e.g., X [101], Instagram [31]) by posting text, photos, or short videos [11, 24]. Regardless of the medium, everyday storytelling is rooted in *immediacy* [69], characterized by minimal curation and editing [95], which enables in-the-moment sharing when the stories remain timely and relevant. Unlike professional forms of storytelling such as novels or films, everyday storytelling is produced by the general public and occupies the lightweight, casual end of the storytelling spectrum [39].

Although *interactive vignettes*—a form of visual storytelling in which the audience, acting as a player, controls a character and participates by interacting with characters and objects (e.g., interactive drama [57, 87], role-playing video games [20, 55])—have gained popularity as an engaging digital medium [25, 42, 83], they have not yet been adopted for everyday storytelling. Interactive vignettes possess two core qualities that can aid immersion of the audience into the narrative: (1) *interactivity* and (2) *first-hand participatory*

experiences; these core qualities *together* give interactive vignettes a potential to enrich everyday storytelling alongside familiar social-media formats such as text, images, and video. However, existing tools for authoring interactive vignettes (e.g., RPGMaker [33], Unity [92], and Roblox Studio [78]) are designed for producing polished works, with development workflows that often span months or years [26]. Interactive vignette authoring often requires authors to curate interactive story scenes [1] and craft narrative branches that diverge into multiple situations [41, 85], which is too technically demanding and time-intensive to accommodate the immediacy required in everyday storytelling.

A formative study (Section 3) with 13 active social media users confirmed the desire to adopt interactive vignettes for everyday storytelling. Participants valued interactive vignettes, especially for everyday stories centered on events, because they can let the audience role-play the author's perspective to unfold events and foster a sense of resonance with the author's personal experience, despite the authoring challenges that hinder immediacy required for everyday storytelling. We further derive three design goals for interactive vignette authoring tools for everyday storytelling: (1) enable a lightweight authoring process that matches the effort of everyday storytelling practices, (2) capture and convey the author's key message in the everyday story, and (3) deliver an engaging viewing experience with situation-aware non-player characters (NPCs) in a branching storyline.

Based on the design goals, we introduce DIARYPLAY (Figure 1), an AI-assisted authoring system for creating interactive vignettes tailored to the lightweight and casual, everyday storytelling context. At a high level, the design of DIARYPLAY explores the balance between human authoring and AI generation in narrative authoring contexts; users can direct the narrative authoring process of interactive vignettes without explicitly specifying all low-level details.

DIARYPLAY takes a natural language text story as input and outputs an interactive vignette, where the player controls the main character (player character; PC) to perform activities by interacting with objects in the environment, and watches other characters (NPCs) perform activities and engage in conversations with NPCs. The authoring system comprises two main components: (1) the *Interactive Vignette Authoring Component*, which supports a lightweight authoring process, through automation and refinement, visually guides the author in specifying the three core elements of an interactive vignette: environment, characters, and events; and (2) the *Interactive Vignette Viewing Component*, which delivers an interactive experience to the player, by enlivening NPCs reacting to player activities and dialogues. The Interactive Vignette Viewing Component relies on the *Controlled Divergence (CD) Module*, an automatic LLM-based module that uses the branch-and-bottleneck [75] storytelling structure to dynamically plan NPC behaviors and allow divergent activities while preserving the author's main messages. With the CD Module, interactive vignettes can branch adaptively without extra author effort on multi-branching.

A preliminary technical evaluation (Section 5) of 16 participants shows that the CD Module generates believable NPC activities based on the given character persona and storyline. To investigate how DIARYPLAY supports everyday storytelling, we conducted a user study with 16 participants (Section 6). The authoring experience met participants' expectations for lightweight effort; they were able

to complete interactive vignette creation within 20 minutes (see resulting examples in Figure 9), and DIARYPLAY provided effective assistance in building interactive vignette elements. Participants valued the moments when players diverged from the main storyline, bringing in tiny yet meaningful pieces that enriched but did not derail their everyday story. Meanwhile, participants described the viewing experience as engaging and immersive, and they were able to understand the author’s main story content despite divergent interactions.

In summary, this paper makes the following contributions:

- DIARYPLAY, an end-to-end AI-assisted interactive vignette authoring system for everyday storytelling;
- the Controlled Divergence (CD) Module, an LLM-powered module that allows divergent player and NPC behaviors within the boundaries defined by the author’s intended story;
- results of a technical evaluation showing that the CD Module generates believable NPC activities based on both character persona and storyline; and
- empirical findings from a user study, demonstrating that DIARYPLAY enables low-effort authoring of interactive vignettes for everyday storytelling while providing engaging viewing experiences and conveying the core story message.

2 Related Work

Our work is related to three main areas of research: (1) interactive vignette and games authoring, (2) building multi-branch narratives, and (3) LLM-powered believable characters in storytelling.

2.1 Interactive Vignette and Games Authoring

In interactive digital storytelling, role-playing [68] is a widely used approach that allows players to control a character and interact with other characters and objects within the story. Role-playing fosters deep immersion and empathy [25, 42], resulting in memorable storytelling experiences [83]. Our work focuses on role-playing through visual storytelling, excluding purely text-based interactive narratives [93, 102]. We adopt the term *interactive vignette* to describe a genre of visual storytelling media that integrates role-playing within an interactive environment (e.g., interactive drama [28, 57, 87, 94] and role-playing games (RPGs) [20, 55, 59]). Our work adopts the framework proposed by Zhao et al. [106], which defines an interactive vignette through three core elements: environment, characters, and events.

Traditionally, creating interactive vignettes has been considered resource-intensive and time-consuming [1, 44, 74, 84]. Authoring tools like Unity [92] and RPGMaker [33] require authors to conceptualize assets, design or source visual representations, manage interactions, and carefully script narrative events. In game design, a growing body of casual creator tools [13] supports lightweight development for non-professional users. Bitsy [15] offers a minimalist 2D tilemap editor for small pixel-art narrative scenes, PuzzleScript [73] provides a domain-specific rule language for rapidly prototyping puzzle games, and Downpour [9] enables collaging text, images, and drawings into simple branching narratives via clickable buttons. While these tools are not designed for everyday storytelling, they reflect a broader trend of empowering the public to create interactive media. To further lower authoring barriers,

research has explored workflows where authors express high-level intent and systems generate interactive artifacts: Germinate [43] and Game-O-Matic [90] use generative models to create games from conceptual inputs, Wevva [72] lets users adjust parameters during gameplay, and RosebudAI [80] uses generative AI to produce game elements from text prompts. However, unlike tools oriented toward game design, DIARYPLAY focuses on everyday storytelling rather than strategy, mechanics, or rules.

Although user-generated storytelling content (e.g., vlogs, blogs) is widespread, interactive vignette creation remains uncommon for everyday storytellers. Recent advances in generative AI, especially LLMs, have streamlined the creation of interactive storytelling elements such as role-play chatbots [45, 88, 106], interactive objects [104], scenes [8], autonomous characters [65], and narrative beats [44]. AI-assisted tools can now generate content directly from natural-language inputs [17], reducing traditional authoring complexity, but may also diminish author agency [19]. Human-in-the-loop [105] systems such as PatchView [12] and FairyTailor [6] address this by pairing AI automation with human modification. Building on this collaborative approach [100], our work uses LLM-driven generation to reduce the authoring burden of interactive vignettes while preserving author agency throughout the process.

2.2 Building Multi-Branch Interactive Narratives

Interactive vignettes foster audience engagement by enabling exploration and influence over narrative progression [44], but they also require narratives to respond dynamically to player interactions while maintaining coherence [4, 76]. Because authors are absent during play, they must predefine structures that anticipate diverse interactions, often through multi-branch storytelling. Kumaran et al. [44] describe two main approaches: representing narratives as story graphs, where states are connected by causal edges and a runtime component manages progression, and enabling the runtime to revise narrative elements dynamically as the story unfolds. In the story-graph approach, tools such as StoryTec [27] require authors to specify branch conditions and actions, and art-E-fact [50] similarly relies on constructing narrative graphs, while AI-assisted systems like GENEVA [47] and Spindle [10] support authors in building these branching structures. In the dynamic-revision approach, StoryVerse [96] combines LLM-driven character simulation with an abstract-act authorial structure to support emergent narratives [3], though it limits authors to specifying high-level conflicts or turning points rather than complete stories. Likewise, Façade [54] provides AI assistance by dynamically selecting and sequencing pre-authored branches (behavior trees) in response to player interactions to maintain a coherent storyline.

However, everyday storytellers, who are often non-expert storytellers, are typically accustomed to single-branch storytelling, and the requirement to conceptualize multi-branch narratives can contradict their natural storytelling habits, making the authoring process challenging [41, 85]. To address this, we aim to develop an implicit approach that transforms single-branch stories into interactive narratives without requiring authors to manually craft multiple branching paths [85]. NarrativeGenie [44] represents an initial attempt. It creates a partially ordered sequence of events based on

an author’s high-level story description and uses an LLM-based pipeline to select and reorder events in response to player interactions. However, in everyday storytelling, event sequence is crucial in preserving the author-intended temporal flow [46]. Hence, we aim to realize branching narrative generation during story runtime that maintains the authored story’s progression while enabling emergent narratives [3] to react to dynamic player interactions.

2.3 LLM-powered Believable Characters in Interactive Digital Storytelling

Virtual characters (e.g., AI agents [65]) are designed to interact with users and simulate real-life behaviors in interactive digital storytelling. To create engaging and immersive experiences, *believability* has been set as a central design goal [65], meaning characters should exhibit lifelike behaviors, make decisions autonomously, and respond naturally to evolving scenarios [7, 89]. A well-defined persona is essential for believable characters [2, 48, 52]. Beyond traditional attribute-based specification, recent work explores conversation-based persona construction in which authors build characters through dialogue with an LLM to better finetune its persona [74]. Commercial tools (e.g., CharacterAI [60]) similarly let users define descriptions, greetings, example dialogues, and personality tags, with iterative refinement through response ratings. Beyond initial character construction, sustaining believable interactions requires a character architecture that maintains persona consistency; the generative agent framework [65] addresses this through memory-based perception, storage, and retrieval to support adaptive yet coherent behavior. Beyond making character behaviors consistent with their persons, our work explores characters that remain believable while *also* aligning with an author’s intended storyline [53, 97].

While existing work primarily focuses on the importance of persona consistency [103], in storytelling, we argue that believable character behaviors should also ensure the narrative’s logical coherence [77]. However, persona consistency and narrative coherence might conflict in a narrative planning [77, 98]. For instance, if a character is portrayed as a person who loves traveling and is set in the storyline to “work hard for a final exam” as a future activity, their behavior becomes less believable if they are shown engaging in a “go on a road trip.” Because it might disrupt the logical flow of the narrative, even if their activity remains consistent with their persona. To address this issue, our work proposes a character activity planning approach that considers both the character’s persona consistency as well as narrative coherence.

3 Formative Study

To understand people’s motivations for authoring and viewing interactive vignettes in everyday storytelling, and to derive design goals for an authoring system to support this practice, we conducted an IRB-approved formative study.

3.1 Procedure

We recruited 13 participants (6 females, 7 males, aged 19 to 35) through online communities within KAIST. To capture potential users of interactive vignette systems, we recruited active social

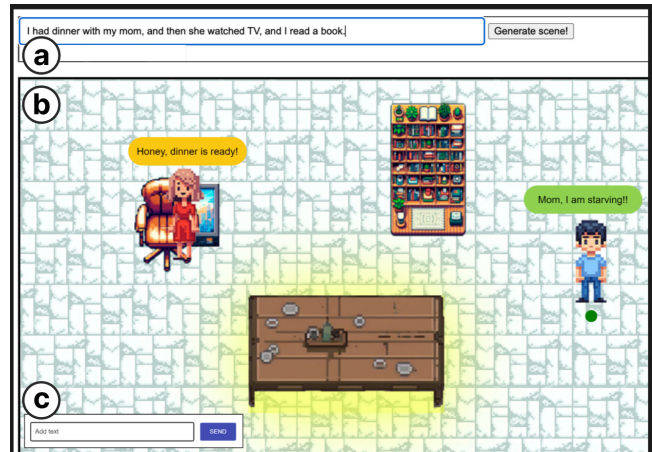


Figure 2: User interface of the research probe, which includes: (a) a text story input box, (b) a display of the automatically generated scene featuring objects mentioned in the story and characters acting on events, and (c) a chat feature.

media users who regularly engage with everyday storytelling, creating and consuming digital storytelling (e.g., social media posts, blogs, vlogs) at least once a week. Participants were from various academic backgrounds and reported various levels of frequency of interactive-vignette consumption (e.g., play RPGs, watch interactive drama), including eight participants who play RPGs at least three times per week. We list detailed participants’ background information in Table 4 at Appendix A. For each participant, two researchers conducted the formative study in person, with one researcher leading the study and the other taking notes. The study lasted roughly 1.5 hours, and each participant received 22,000 KRW (≈ 17 USD) as compensation.

After a brief introduction, we conducted a semi-structured interview. The first half of the interview aimed to explore the motivations for authoring and viewing interactive vignettes in everyday storytelling. We introduced interactive vignettes alongside four common everyday storytelling modalities in social media (single image, sequential images, text, and video), presented in random order with balanced descriptions. We asked the participant to rank five modalities from most to least preferred with 1 being best, both as an author and as an audience, allowing us to assess their intrinsic motivation to create and engage with interactive vignettes in everyday storytelling. From the author’s perspective, we inquired about their motivations for authoring and their expected outcomes for an interactive vignette representing everyday stories, assuming they would have adequate authoring supports.

In the second half of the interview, we shifted focus to the participants’ expectations for an interactive vignette authoring system. The goal of the second half of the formative study was to gather more specific insights about what an ideal output of an interactive vignette authoring system would look like for everyday storytelling, as well as to understand where the authoring system could assist authors in the authoring process. To ground the discussion in concrete examples, we asked participants to interact with a research

probe (Figure 2) that we developed. This probe uses an LLM to automatically convert a short text story into a scene with objects and characters mentioned in the story. We adopted the three core elements of interactive vignettes (environment, characters, and events) identified in the literature [106]. Additionally, we drew inspiration from conventional tile-based 2D RPGs (e.g., Pokémon [22], Chrono Trigger [86]) and Generative Agents [65], which demonstrate object-centric character activities and conversational interaction. Also, we followed the recent trend of storytelling authoring tools (e.g., NarrativePlay [106], SceneCraft [45]) that use natural language story as an intuitive input modality. In designing the probe, we intentionally kept the affordances minimal yet expressive, so that participants could articulate their expectations, preferences, and the missing features for everyday storytelling.

For the analysis of the collected data, the first and second authors performed open coding [36] on the interview transcripts and observation notes, identifying themes and consolidating them through iterative discussions among the whole research team. We provide further descriptions of the research probe, study protocol, and interview questions in the Supplementary Material.

3.2 Results

We present participants’ motivations for authoring interactive vignettes in everyday storytelling, their expectations for the authoring process, and their desired outcomes and viewing experiences.

3.2.1 Motivations for Adopting Interactive Vignettes in Everyday Storytelling.

Participants expressed a desire to both author and view interactive vignettes in everyday storytelling. Compared to other modalities, interactive vignettes had a higher average ranking from both the authoring ($\mu = 1.77$) and viewing ($\mu = 1.61$) perspectives. From the author’s side, participants valued the role-playing aspect, which allowed them to “actively invite the audience to experience my (the author’s) perspective and deeply resonate with my everyday story” (P7). From the audience’s side, participants highlighted the value of “proactive participation in the events beyond traditional passive consumption of an everyday story” (P9). A Friedman’s test showed significant differences in modality preference rankings for both authoring ($\chi^2 = 15.38, p = 0.004$) and viewing ($\chi^2 = 28.96, p < 0.001$). Nemenyi post-hoc tests (Table 1a and Table 1b) indicate that from both authoring and viewing perspectives, interactive vignettes ranked significantly better than text (authoring: $\mu = 4.00, p < 0.01$; viewing: $\mu = 4.39, p < 0.01$) and single images (authoring: $\mu = 3.54, p = 0.04$; viewing: $\mu = 3.79, p < 0.01$).

Regarding suitable formats for various story types, participants noted that interactive vignette is well suited for everyday stories with clear event sequences (P1, 4, 8, 9)—“move characters from one event object to another, feels very engaging.” (P4)—but may be less expressive for stories centered on emotions or inner reflection (P6, 7, 11). Instead, P14 suggested the potential benefits of a text story for such stories, since text stories allow “putting more details about inner emotions, that are hard to express through visual language and can only be conveyed through words.”

3.2.2 Expectations for a Lightweight Authoring Process and Automation Support.

All participants emphasized they expected the process to remain lightweight. As participants self-reported, they usually

Table 1: Pairwise p-values between five modalities were calculated using Nemenyi’s test. The cells for significant pairs ($p < 0.05$) are bolded. The average ranking for each modality is shown after the modality names. We denote the five modalities as interactive vignette (IV), multiple images (MI), single image (SI), video (VD), and text (TX).

(a) Authoring side					
	IV ($\mu=1.77$)	MI ($\mu=3.08$)	SI ($\mu=3.54$)	VD ($\mu=2.62$)	TX ($\mu=4.00$)
IV	-	0.22	0.04	0.65	< 0.01
MI	-	-	0.95	0.95	0.57
SI	-	-	-	0.57	0.95
VD	-	-	-	-	0.17
TX	-	-	-	-	-

(b) Viewing side					
	IV ($\mu=1.61$)	MI ($\mu=2.96$)	SI ($\mu=3.79$)	VD ($\mu=2.25$)	TX ($\mu=4.39$)
IV	-	0.15	< 0.01	0.82	< 0.01
MI	-	-	0.64	0.75	0.12
SI	-	-	-	0.08	0.85
VD	-	-	-	-	< 0.01
TX	-	-	-	-	-

invested about 20 minutes in creating content for everyday storytelling, setting their expectation for a similar duration in interactive vignette creation. Participants considered text-story input intuitive and easy to use. At the same time, participants perceived the complexity of authoring interactive vignettes and emphasized the need for a “low-effort” and “not time-consuming” process (e.g., creating various elements of an interactive vignette from scratch is cumbersome), reaffirming the importance of immediacy in everyday storytelling. As P8 noted, “Everyday storytelling should be a lightweight task, not too time-consuming or laborious.” They therefore appreciated the probe’s automation based on natural language text-story input, even though it was not perfect. As P4 noted, “changing things is easier than building from scratch”, which reflects the “no blank canvas” concept [13].

3.2.3 Expectations for the Resulting Interactive Vignettes.

Participants believed the resulting interactive vignette should convey the main message of their story—who did what, where—consistent with the “environment, characters, and events” framework [106] and maintains “one beginning, one ending” (P7). For the environment, participants preferred an object-rich design (e.g., “a sandbox environment like RPGs,” P9) rather than sparse, randomly positioned objects tied only to specified activities. For characters, participants valued conversational features and expected to represent character personas (e.g., mood, personality, social relationships). In contrast, visual details such as clothing and facial features were “not very important for storytelling” (P2). For events, participants expected the outcome to reflect the event sequence described in their text-story input and to show characters moving through the scene between objects, which felt “vivid and immersive” (P3). In sum, participants suggested the authoring system should add object-rich environments with better realism, enable persona building in character construction while retaining the conversational feature, and keep character mobility within the environment for event transitions.

3.2.4 Expectations for Player Divergence in Viewing Experience. Participants emphasized that an engaging viewing experience should involve flexible interaction with objects in the environment and conversations with NPCs. When authoring, participants wanted their input story to serve as a “*central skeleton with key events as checkpoints*” (P5), allowing players to go beyond the predefined key events. As P10 noted, forcing players to strictly follow a fixed author-defined storyline would feel “*too boring, like finishing a homework assignment*.” Participants also expected NPCs to respond with adaptive actions when players diverged from the main storyline, rather than remaining idle, to sustain immersion and narrative coherence.

3.3 Design Goals

Based on the findings, we outline three design goals for an interactive vignette authoring system for everyday storytelling.

[DG1] Enable a lightweight authoring process that matches the effort of everyday storytelling practices. The system should support single-branch natural-language story input and provide system automation for generating narrative elements.

[DG2] Capture and convey the author’s key message in the everyday story. The system should enable authors to concretize the core elements (environment, characters, and events) of the interactive vignette and ensure that the generated interactive vignette conveys them.

[DG3] Deliver an engaging viewing experience with situation-aware NPCs in a branching storyline. The system should support divergent player interactions and provide adaptive branching narratives. To preserve the author’s intended events (DG2), the system should use a narrative structure with one beginning and one ending, guiding players through events in the order defined by the author.

4 The DIARYPLAY System

Based on the findings and design goals (Section 3), we present DIARYPLAY, an AI-assisted interactive vignette authoring system for everyday storytelling. The system comprises two components (Figure 3): (1) the *Interactive Vignette Authoring Component*, which includes the *Authoring Interface* and the *Interactive Vignette Elements Extractor*, and (2) the *Interactive Vignette Viewing Component*, which includes the *Viewing Interface* and the *Controlled Divergence Module*. To support lightweight authoring (DG1) and ensure the capture of the author’s core story message (DG2), the Interactive Vignette Authoring Component guides everyday storytellers in completing and refining interactive vignette elements by surfacing gaps between the natural-language story input and the information required for an interactive vignette. To enable adaptive player interactions (DG3) without requiring authors to manually construct branching paths (DG1), the Interactive Vignette Viewing Component automatically delivers branching interactive narratives based on the authored elements. We adopt a branch-and-bottle narrative structure (Figure 4) so that divergent player interactions (DG3) converge back into the author’s predefined key events, thereby preserving the author’s intended story message (DG2).

4.1 Interactive Vignette Elements Specification

Based on previous literature [106] and the formative study, we use three core elements to represent interactive vignettes: *environment*, *characters*, and *events*. To enable lightweight authoring (DG1) while still capturing the main story message (DG2), we designed a minimal but sufficient specification for everyday storytelling regarding each of the three core elements. We determined the input fields for each of the elements by referring to previous work and systems, feedback from the formative study, and iterative pilot tests to ensure a reasonably comprehensive representation. We note that the purpose of the specification is not to capture comprehensive properties of the three elements that would generalize beyond the everyday story (e.g., a full-scale persona), but to specifically capture the required information to play out the everyday story as a means to minimize authoring effort (DG1).

- **Environment:** The visual setup of the interactive scene in which the story unfolds. For our system, we use a 25x35 2D grid map with walled spaces (rooms with entrances) to represent a diversity of locations (e.g., office, home, grocery store, theater, etc). The environment includes the layout of the rooms in the scene as well as the placement of the objects. Each object in the scene is represented as its location, scale, actions characters can perform with the object, and interaction zones (i.e., where characters have to be spatially when interacting with the object (e.g., on top, in front, next to)).
- **Characters:** Persona and visuals of the player character (PC) and the non-player characters (NPCs). Based on the character generation forms in existing AI character constructions [60, 65, 74], the fields for each of the characters include the name, age, personality, social role (e.g., relationship to other characters, occupation), current mood, language style, examples of conversation snippets for the persona, and character sprites for the visuals.
- **Events:** An ordered sequence of *key events*, each of which is a group of simultaneous character activities (e.g., Kelly cooks dinner while Julie helps with cooking, Event 1 of Figure 3). Each character activity is represented as a (character, action, object) tuple, which specifies “*who (character) did what (action) while interacting with what (object)*”. The order of the key events indicates the chronological order of the events (before/after). We note that not all characters need to have an assigned activity for each key event if not specified in the input story.

The information available in the specification allows the Interactive Vignette Viewing Component to automatically generate a branch-and-bottleneck narrative structure by enlivening the NPCs who interact with the environment in a way that is consistent with the NPCs’ personas and the key events.

4.2 Usage Scenario

We illustrate DIARYPLAY through a usage scenario that walks through both the authoring and viewing scenarios. Kelly, who enjoys sharing her daily moments with friends on social media, uses DIARYPLAY to create an interactive vignette and share it with her friend Bob, who experiences it as a player.

4.2.1 Authoring Scenario & Interface. To author an interactive vignette using DIARYPLAY, Kelly proceeds through the following steps in the *Authoring Interface* (Figure 5a).

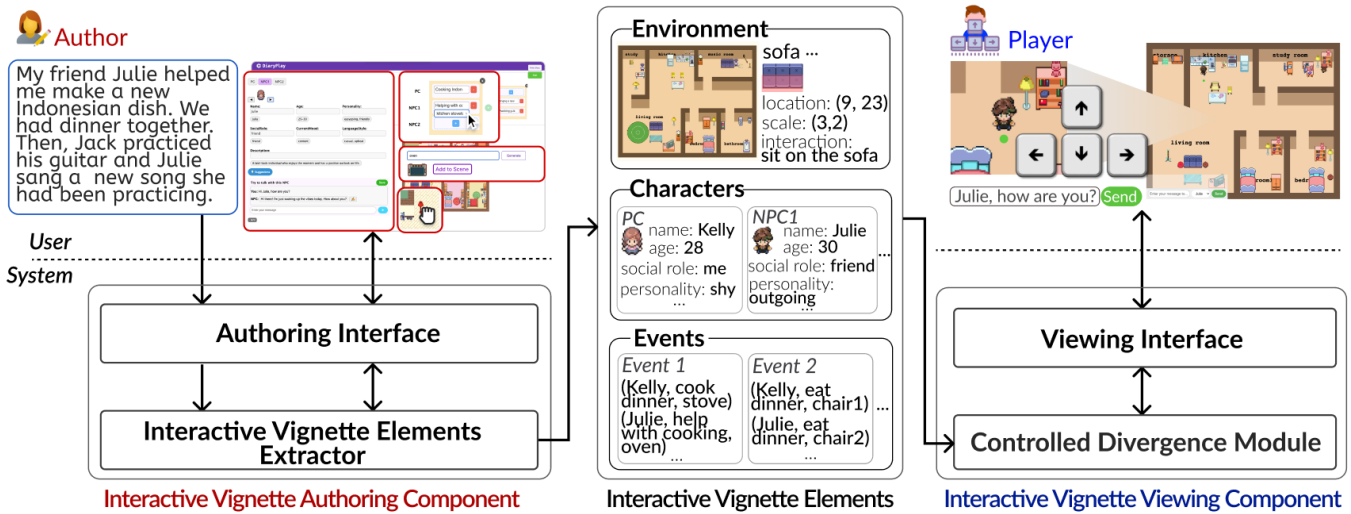


Figure 3: DIARYPLAY system overview. The Interactive Vignette Authoring Component takes the author’s text story as input, extracts core elements (environment, characters, and events), and guides the author in refining them into structured interactive vignette elements. The Interactive Vignette Viewing Component enlivens NPCs and allows the player to control the PC, enabling an interactive narrative experience grounded in the authored storyline.




Figure 4: Illustration of the branch-and-bottleneck structure in interactive storytelling. Solid black circles represent key events that all players encounter, ensuring alignment with the author-intended storyline. Hollow circles represent divergent events, which offer optional or personalized paths that allow for variation in viewing experience. Despite branching paths, the narrative converges at key events.

Step 1. Input a story. Kelly begins by entering an everyday story¹ into the story input box: *My friend Julie helped me make a new Indonesian dish. We had dinner together. Then, Jack practiced his guitar and Julie sang a new song she had been practicing.* Once the system processes the input, Kelly sees three panels, displaying each of the initial extraction results of three interactive vignette elements (Figure 6a).

Step 2. Label rooms. Kelly decides to begin with the Environment panel, where she sees a home map layout with several empty pre-labeled rooms (e.g., “bedroom,” “kitchen”; Figure 6b). She revises the “dining” to “living room” to better align with her intended setup, then clicks the “Confirm the rooms” button to continue.

Step 3. Place objects. Soon, she sees the rooms automatically populated by the system (Figure 6c), including objects relevant to each room (e.g., a bed for each bedroom) and objects required for the

events in the story (e.g., a guitar for Jack to practice). To customize the environment, Kelly uses mouse drag-and-drop to reposition objects. Noticing that the living room feels crowded, she right-clicks to remove a carpet. She notices a chair is missing and then decides to add a dining chair in the kitchen, so she types “dining chair” in the input box above the map, and clicks the “Generate” button, then places the new chair in her desired location.

Step 4. Complete the character profiles. Next, Kelly moves to the Character Panel. She notices three tabs (Figure 6d), each representing a character from her input story: the Player Character (PC), which will be controlled by the player; and two NPCs, her friends Julie and Jack. Each character tab includes an avatar selection and several blank fields for defining their personas. Kelly customizes the avatars by cycling through options using arrow buttons. When filling in the persona fields, she struggles to articulate Jack’s personality, so she uses the “Try to talk with this NPC” feature to simulate a real-life conversation she had with him. After reviewing the system-generated response, she directly changed the system-generated dialogue (*Hey Kelly, don’t worry too much!* 🙌) to better reflect how Jack would actually speak in the real world. Once satisfied with the conversation sample, she clicks the “Suggestions” button and receives a suggestion bubble beneath the personality field. Agreeing with the suggestion (“supportive”), she clicks the bubble to automatically fill in the blank field.

Step 5. Refine the events with object assignment. After completing the Characters panel, Kelly navigates to the Events panel, where she finds a visual timeline of key events (Figure 6e), each containing character activities. She carefully reviews each event to verify that the system’s extracted activities align with her input story. Kelly noticed that an event generated from “we had dinner” included only Julie and her, but it should have included Jack too. To

¹The example is collected from a pilot study.

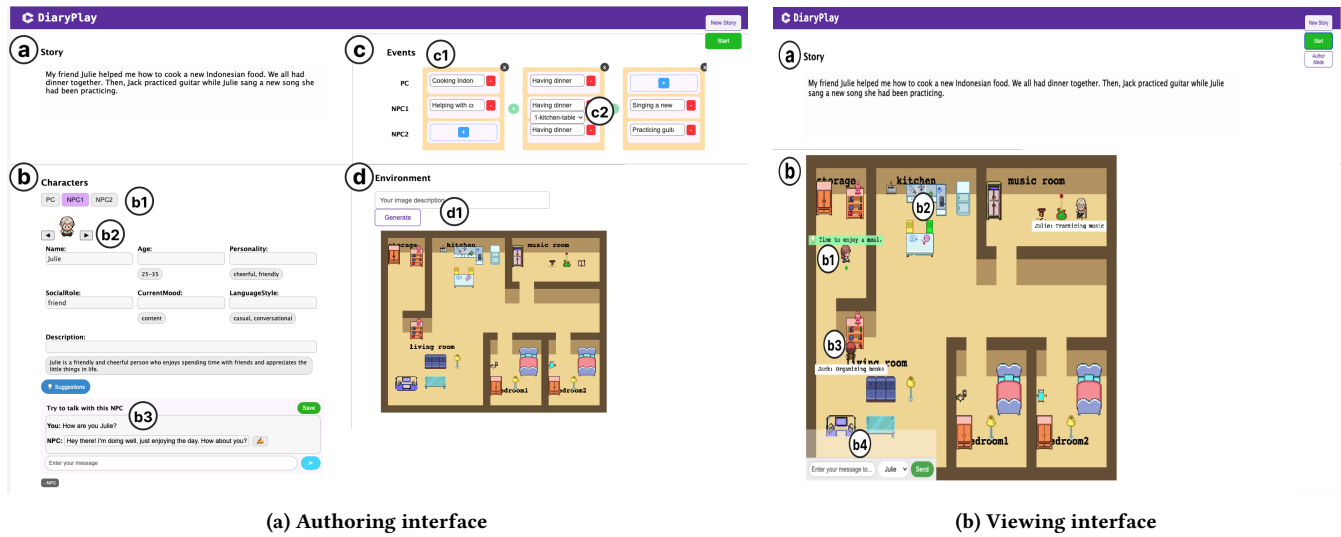


Figure 5: Authoring interface, including (a1) Text story input, (a2) Characters panel, (a3) Events panel, and (a4) Environment panel; and Viewing interface, including (b1) Text story caption and (b2) Interactive vignette.

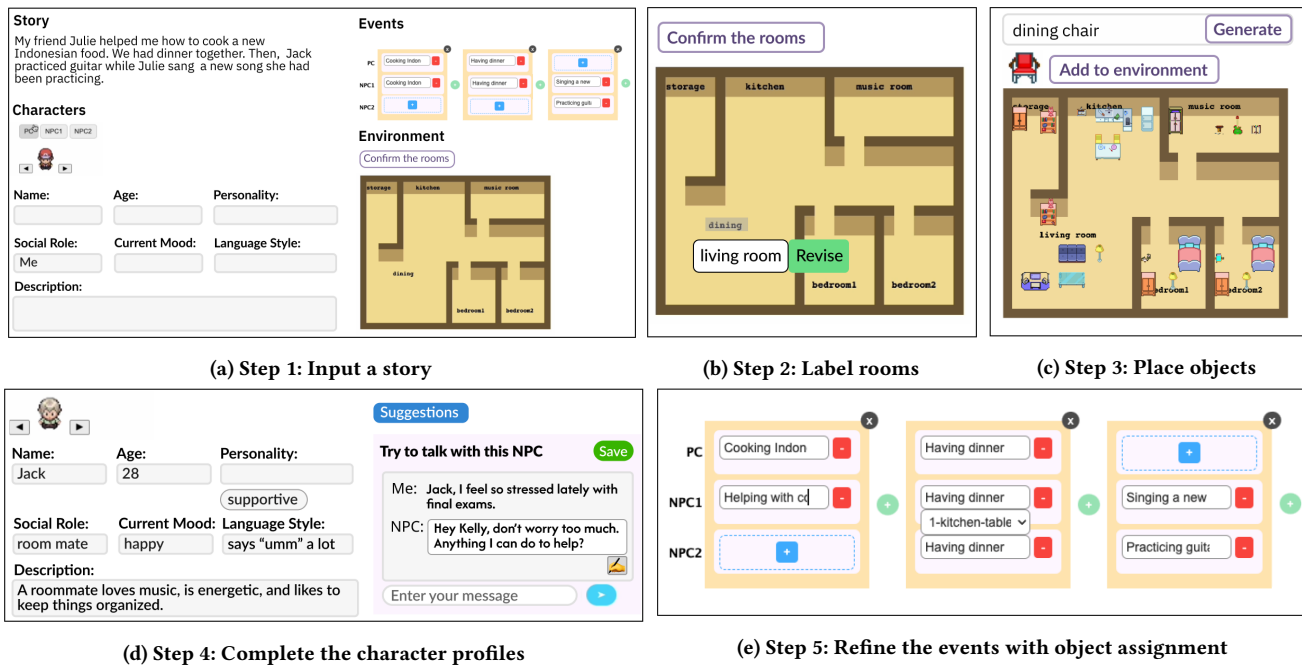


Figure 6: Steps in the Authoring Scenario.

correct this, she clicks the add button (+) and types “having dinner” to add it to Jack’s activities in the events panel. After making this revision, Kelly feels that the events are now in-line with her intended story, with character activities arranged chronologically and simultaneous activities grouped within the same event. Next, to review and edit the target objects for each activity, she clicks on each activity and sees the system’s automatic assignments and opens the drop-down menu to select an object from the list.

Satisfied with the elements of the three panels, Kelly clicks the “Start” button to preview her interactive vignette (Figure 5b). If she wants to refine it further, she can click “Author mode” to continue making adjustments.

4.2.2 **Viewing Scenario & Interface.** After Kelly posts the interactive vignette on social media, Bob plays it through interacting with *Viewing Interface* (Figure 5b).



Figure 7: Steps in the Viewing Scenario.

Step 1. Read the caption. Bob starts by reading the text story as a caption (Figure 7a) to get an overview of the characters and events. He realizes that he will role-play as Kelly (the PC) and interact with two other characters, Julie and Jack (NPCs).

Step 2. Engage with the first key event. As the interactive vignette loads, Bob notices the kitchen stove top glowing (Figure 7b), signaling that an activity is about to take place at the stove. To trigger the activity, he uses the arrow keys to move the PC toward it. Upon approaching, the activity starts. Bob sees the glow disappear, and a speech bubble appears under the PC's avatar to show the PC's activity: "cooking dinner." Bob then observes Julie moving in the kitchen, and her speech bubble reads: "helping with cooking." Although the text story does not specify Jack's activity during the cooking, Bob notices that Jack is sitting on the sofa and watching TV rather than remaining idle, giving Bob a sense of character liveliness. As the cooking event continues, Bob engages with the interactive vignette by role-playing as Kelly and talking with Julie. He types into the message box to ask Julie, "How much spice should I add?" Julie responds, "Just a pinch or to taste!"

Step 3. Diverge from the key events. After the cooking event is done, Bob sees the chair glowing, signaling the transition to the next key event — having dinner. However, instead of moving the PC to the glowing chair (Figure 7c) immediately, Bob decides to explore the home environment further. He moves the PC to the storage room and triggers the activity "cleaning the bookshelf" at the bookshelf. Interestingly, Bob realizes Jack and Julie are not abandoning the PC to have dinner without Kelly. Instead, Bob observes that Jack and Julie are doing other activities. Jack is "organizing the wardrobe", and Julie is "practicing music." Meanwhile, Bob receives a subtle inner voice bubble reading: "Time to enjoy a meal!" Trying to diverge from eating dinner, Bob types to Jack: "I want to skip dinner." Jack responds, "Dinner is important. Let's have dinner together," guiding the PC toward the next key event.

Step 4. Return to the main storyline. Following the inner voice and Jack's response, Bob decides to move the PC to the glowing dining chair. After the PC arrives at the dining chair, Bob sees Jack and Julie join the PC for dinner (Figure 7d).

Step 5. Complete the interactive vignette experience. After the “having dinner” event, Bob notices Jack and Julie moving to the music room to practice guitar and singing (Figure 7e). With no more glowing objects in the environment, Bob realizes he is free to move the PC, interacting with the characters and objects in the environment. After Jack and Julie complete their guitar practice and singing, Bob sees the end screen, marking the conclusion of the experience.

4.3 System Implementation

To support the usage scenario described above, we implement DIARYPLAY with two modules: (1) the Interactive Vignette Elements Extractor supporting the authoring workflow, and (2) the Controlled Divergence Module running during the viewing experience.

4.3.1 Interactive Vignette Elements Extractor. The Interactive Vignette Elements Extractor takes the author’s input story and initially extracts three elements — environment, characters, and events — and presents them in the corresponding panels on the Authoring Interface, where the author can review, complete, and modify them with system guidance.

Environment Building: Systems for interactive storytelling often require authors to construct environments for characters to act within. To avoid overwhelming authors with the heavy workload of building environments from scratch (DG1), the Interactive Vignette Elements Extractor automatically generates the environment layout and places objects based on the full story input, with each step followed by user confirmation and modification. This story-driven environment building echoes Sleep is Death [49], where one player builds the environment for another’s story, whereas here the system assists in generating the environment for the author’s story.

To first define the environment layout, the Interactive Vignette Elements Extractor begins by using the LLM to pair the input story with the tags (e.g., residential, retail, office) on one of the layouts (environment templates with a fixed number of rooms) designed based on the urban space classification by Pissourios et al. [71]. Then, it uses the LLM to label each of the rooms (e.g., kitchen, bedroom) based on its functionality to support the story input, and asks the author for review and revision.

Once the author confirms the layout and the labels for each of the rooms, the Interactive Vignette Elements Extractor then automatically calculates and populates each of the rooms with *necessary objects* (either event-related (e.g., guitar for Jack’s practice) or environment-related (e.g., bed in a bedroom)) using a greedy algorithm. If there is still room for additional objects, the extractor generates *decorative objects* that enhance realism (e.g., a lamp in a bedroom). The positioning of the objects is based on the LLM’s semantic reasoning of object-object relationships (e.g., “chairs” close to “tables”), and object-room relationships (e.g., “fridges” next to “walls”), as well as the path-finding algorithm to ensure the presence of paths to the objects. Once the population of the rooms is complete, the authors can add, remove, or relocate objects. Although there is no fixed limit on room objects, the path-finding algorithm stops additions when the room becomes too crowded for navigable paths.

For each of the objects, the system assigns possible actions based on the story and appropriate trigger zone type — where the character sprite needs to be with respect to the object to trigger an

interaction — by prompting the LLM with the object’s name to reason how people typically interact with it: *on* (e.g., sleeping on a bed), *partial* (e.g., sitting on the seat portion of a sofa), *around* (e.g., sitting around a table), and *directional* (e.g., opening a fridge from the front side).

Characters Construction: Based on the interactive vignette elements specification about the characters, the extractor utilizes the LLM to analyze the characters mentioned in the story. It identifies the main character, first-person narrator (e.g., “I”) as the PC and classifies all other characters as NPCs. For each character, the system extracts explicitly stated persona attributes such as names and social roles, while leaving unspecified traits blank for the author to complete, with system suggestions available to ease authoring. Although the LLM is capable of inferring implicit attributes, such as a character’s mood, we intentionally avoid auto-filling these to prevent over-guiding the author.

The conversation-simulation feature aligns with prior work in computational character authoring, which uses LLMs to generate dialogue based on character persona traits [60] and to infer character persona traits based on given dialogues [74]. The system prompts an LLM to generate NPC replies based on the character’s current persona traits, conversation examples, and the story context. Authors can edit these responses to better reflect how the real person would speak. To recommend persona traits, the system prompts an LLM to compare the dialogue against current persona traits: it flags traits contradicted by the conversation, revises them, and proposes new traits implied by the dialogue. The recommended list includes retained, revised, and newly inferred traits, which authors can further modify. By leveraging the interdependence between dialogue and persona traits, DIARYPLAY supports both explicit and implicit character creation through iterative refinement.

Events Scheduling: Based on the interactive vignette elements specification about the events, the extractor first prompts the LLM with the characters extracted during character construction and the input story to list the actions performed by each character, matching them to the most appropriate objects in the environment, thereby creating the (character, action, object) tuple. Next, the extractor prompts the LLM to group simultaneous activities by multiple characters into a single key event. Finally, the system asks the LLM to organize these key events into a time-ordered sequence.

4.3.2 Controlled Divergence Module. During the interactive vignette play time, the Controlled Divergence (CD) Module automatically transforms the single-branch sequence of events in the specification into a branch-and-bottleneck narrative structure reacting to the player interactions. The purpose of the CD Module is to allow both the PC and the NPCs to freely take *divergent activities* (DG3), while being subtly guided to stay within the *controlled boundaries* defined by the storyline from the author (DG2). From the interactive vignette author’s perspective, the CD Module offloads the manual effort of crafting a multi-branch narrative (DG1); from the player’s perspective, the CD Module provides player agency within the viewing experience (DG3).

We considered three principles when designing the CD Module. First, to ensure DG3, the CD Module must provide *real-time* responses to ensure an immersive and coherent narrative experience.

Second, to support DG2, the CD Module should provide *subtle PC guidance* toward the key events when they diverge, to preserve the author-intended storyline. Third, to support DG2, the CD Module should ensure *believable NPCs* that behave consistently with their personas and the overarching story logic. Based on these principles, the CD Module plans NPC activities and reacts to PC activities in a two-stage loop (Figure 8):

Stage 1: Plan NPC activities. To ensure real-time responsiveness despite the computational delay from activity generation, the CD Module plans NPC activities in advance, triggering planning for the next activity as soon as the current activity begins. The design of the CD Module takes inspiration from how the human brain copes with the delay between the sensing of the environment and the processing of the information by generating multiple possible action plans and then pruning them according to incoming information [34].

Specifically, the CD Module generates two *next* activity plans for each NPC using the LLM based on the NPC’s *persona* and the *storyline*, which includes all past activities, the characters’ ongoing activities and the activities in next key events (Figure 8 top): one for when the PC takes action according to the key event (may be defined in the specification; possibility A), and another for when the PC engages in a divergent activity (possibility B).

Stage 2: React to PC activities. Regarding the NPC activities, the CD Module continuously monitors PC activities and takes the activity A or B based on whether the PC took the activity defined in the next key event (Figure 8 bottom). Once the NPC begins taking the planned action, the CD Module returns to Stage 1 to plan the next activity.

The CD Module also detects the PC’s deviation from the key events to guide the player back to the flow of the story through two mechanisms: *inner voice* and *chat-based guidance from the NPCs*. The inner voice is triggered when the PC does not take the activity defined in the next key event; the CD Module prompts the LLM to write out what the next activity is in a green thought bubble (e.g., Figure 7c; “Time to enjoy a meal”). When the PC’s chat shows an intention to deviate too far from the key events, the CD Module uses the NPCs’ chats to guide the PC back. For example, if the PC attempts to skip a key event (e.g., “Can I skip dinner?”), the NPC would not agree to the PC but rather provide a response suggesting returning to the author-intended storyline, “Dinner is important. Let’s enjoy it together.” Instead of forcing players back to the key event, the system uses inner-voice cues and dialogues to nudge players [81]. To preserve immersion, we allow divergent activities to continue as long as needed without a fixed limit.

4.4 Implementation Notes

We developed DIARYPLAY with Phaser [70]. The visual assets are sourced from online open resources [32] and generated by DALL-E 2 [62]. We used OpenAI’s gpt-4o-2024-11-20 [63]. We include LLM prompts in the Supplementary Material. To manage computational costs, the current version of DIARYPLAY supports a maximum of three characters. To prevent harmful conversations, the CD Module monitors dialogues and withholds responses if they violate OpenAI’s usage policies.

5 Preliminary Technical Evaluation

The NPC activities generated by the CD Module are central to delivering an immersive and interactive experience for players. In this preliminary technical evaluation, we focused on the CD Module’s ability to generate believable NPC activities, since believability is widely recognized as a crucial aspect of narrative quality [14, 23, 48]. We performed an IRB-approved preliminary technical evaluation with five conditions:

- **CD Module Condition [CD]:** We use the CD Module of our system to generate NPC activities based on persona and storyline properties.
- **Baseline Condition [BL]:** We use a simple baseline module that randomly selects an object in the environment and prompts the LLM to assign an NPC activity.
- **Human Author Condition [HA]:** We asked the original author of each of the interactive vignettes used in the study to directly fill in the NPC activities. This condition serves as an optimal solution to the activity generation problem.
- **Persona Only Condition [PO]:** We generate NPC activities with our CD Module with only the persona and without the storyline property.
- **Storyline Only Condition [SO]:** We generate NPC activities with our CD Module with only the storyline and without the persona property.

With these conditions, we consider two hypotheses:

[H1]: The CD Module is capable of generating NPC activities more believable than the baseline and on par with the human authors themselves.

[H2]: Both the persona and the storyline contribute to generating more believable NPC activities.

We note that the preliminary technical evaluation was performed with an earlier version of the system that only included limited environment editing features. However, the system included complete support for character and event extraction, which are highly relevant for the preliminary technical evaluation.

5.1 Method

To first obtain the interactive vignettes that we could generate the NPC activities on for each of the five conditions, we recruited 16 participants (7 females and 9 males, aged 19 to 30) within KAIST through recruitment posts. Each participant authored an interactive vignette with DIARYPLAY. Next, another participant interacted with the interactive vignette as the player; everyone triggered the CD Module to either fill in unspecified activities or divergent activities (on average three times). We then showed the character-activity table (Table 2) to the original author of each interactive vignette with the CD Module’s generated NPC activities redacted and collected the author’s intended NPC activities for the HA Condition. For the two ablation conditions (PO, SO), we used the appropriate variants of the CD Module to generate NPC activities.

Next, to compare the conditions, we recruited a disjoint set of 16 evaluators (6 females and 10 males, aged 18 to 33) within KAIST through recruitment posts. After introducing the character-activity table, we provided evaluators with the character attributes and input story as supporting material, along with the character-activity tables generated under five different conditions. We then asked

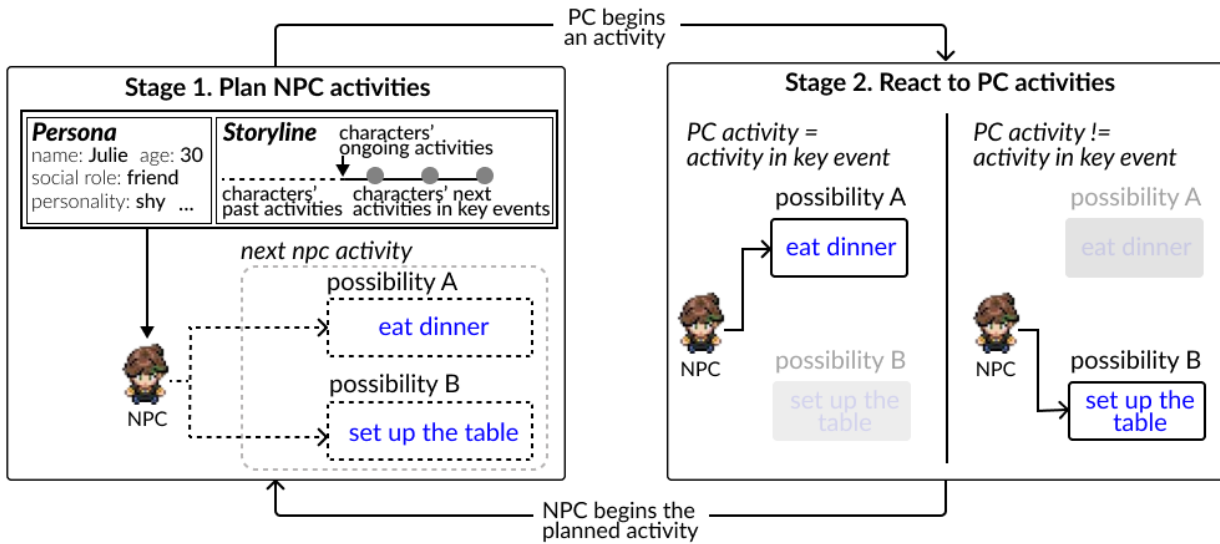


Figure 8: The CD Module plans NPC activities and reacts to PC activities in a two-stage loop.

Table 2: A character/activity table prints the activities of all characters in a temporal order. NPC activities generated during the interactive vignette viewing period, or added later by human authors, are highlighted in green.

Events	PC (Me)	NPC1 (Name)	NPC2 (Name)
1	(sleep, bed)	(cook dinner, oven)	(watch TV, sofa)
2	(eat dinner, chair1)	(eat dinner, chair2)	(eat dinner, chair3)
...
n	(shower, bathroom)	(painting, art desk)	(do laundry, dryer)

them to rank the believability of the NPC activities (highlighted in green in Table 2) across these five conditions (1 is the best, 5 is the worst). Note that evaluators saw activity tables generated in all 5 conditions in random order without notations. We then asked the evaluators to leave any free-form comments on each of the rankings. We did not let technical evaluators experience interactive vignettes to minimize confounding distractions (e.g., visuals) beyond text-based NPC activities. Throughout the evaluation, we did not describe how the activities were generated and randomized the order in which the five conditions were shown to avoid potential biases. Each evaluator ranked the NPC activities from 12 sets of conditions; we collected a total of 192 sets of rankings.

Each of the participants and the evaluators received 25000 KRW (≈ 19 USD) for 2 hours of their time. We include detailed preliminary technical evaluation materials in the Supplementary Material.

5.2 Results

The average rankings of the character activities generated for each condition were 2.37 for CD, 2.58 for HA, 2.74 for PO, 3.31 for SO, and 4.00 for BL.

A Friedman’s test shows that there exist significant differences in the ranking of the believability of the conditions ($\chi^2 = 133.78, p <$

Table 3: Pairwise p-values between five conditions were calculated using Nemenyi’s test. The cells for significant pairs ($p < 0.05$) are bolded. The average ranking for each condition is shown after the condition names.

	CD ($\mu=2.37$)	BL ($\mu=4.00$)	HA ($\mu=2.58$)	PO ($\mu=2.74$)	SO ($\mu=3.31$)
CD	-	< 0.01	0.71	0.14	< 0.01
BL	-	-	< 0.01	< 0.01	< 0.01
HA	-	-	-	0.84	< 0.01
PO	-	-	-	-	< 0.01
SO	-	-	-	-	-

0.001). To understand where the pairwise differences are, we apply the Nemenyi post-hoc test (Table 3); we analyze each of the differences further along each of the hypotheses.

5.2.1 Assessing H1. The character activities generated by the CD Module ($\mu = 2.37$) ranked significantly better than those generated in the baseline condition ($\mu = 4.00; p < 0.01$). As E3 commented, character activities generated by the CD Module were “*more temporally logical and relevant to previous and later activities*” and “*captured established characters well,*” but those generated in the baseline conditions were “*quite random, did not make much sense.*”

We did not find a significant difference between the ranking of character activities generated by the CD Module ($\mu = 2.37$) and human authors ($\mu = 2.58; p = 0.71$). This result indicates that the CD Module generates character activities on par with human authors regarding the believability. In sum, the results support [H1]: the CD Module generated more believable NPC activities than the baseline and was on par with the human authors.

5.2.2 Assessing H2. The character activities generated in the PO condition ($\mu = 2.74$) ranked significantly better than those generated in the baseline condition ($\mu = 4.00; p < 0.01$). This result suggests that persona contributes to believable character activity generation. In addition, the character activities generated in the

SO condition ($\mu = 3.31$) ranked significantly better than those generated in the baseline condition ($\mu = 4.00$; $p < 0.01$). This result suggests that the storyline contributes to believable character activity generation.

The character activities generated by the CD Module ($\mu = 2.37$) ranked significantly better than those generated in the SO condition ($\mu = 3.31$; $p < 0.01$), suggesting that combining persona with storyline leads to more believable activity generation than using storyline alone. The character activities generated by the CD Module ($\mu = 2.37$) ranked better than those generated in the PO condition ($\mu = 2.74$; $p = 0.14$) on average, although we did not see a significant difference between PO and CD. In sum, the results support [H2]: both the persona and the storyline contributed to generating more believable NPC activities.

6 User Study

We conducted an IRB-approved user study to understand the authoring and viewing experiences of interactive vignettes using DIARYPLAY for everyday storytelling. Specifically, we gathered qualitative feedback about the system around the research questions (RQs):

[RQ1]: How well does DIARYPLAY assist interactive vignette authoring for everyday storytelling?

[RQ2]: How well does DIARYPLAY capture and convey the author's main story messages?

[RQ3]: How is the viewing experience delivered by the interactive vignette generated with DIARYPLAY?

6.1 Participants

We recruited 16 participants (3 females and 13 males, aged 18 to 26) through recruitment posts within KAIST (See detailed participants' background information in Table 5 at Appendix A). Participants possessed various academic backgrounds, and were active social media users, creating and consuming everyday storytelling content (e.g., image posts, journaling, text updates, personal vlogs) more than once per week. Many participants reported experience with AI-assisted authoring tools for text and image generation.

In each session, we randomly paired two participants and had them alternate roles as author and player: the player shared their takeaway messages, and the author responded on whether they matched the core story message (who did what, where), an approach used in evaluating authoring tools [51]. Each user study session lasted around 1.5 hours, and each participant received 30000 KRW (≈ 21 USD) as compensation for their time.

6.2 Procedure

We began each in-person user study session by introducing the study's purpose, key terms, and a tutorial on using DIARYPLAY (about 5 minutes).

Following the introduction, we let each participant create an interactive vignette based on their personal daily moment, single-location story with DIARYPLAY on the provided device independently. Throughout the authoring process, we asked the participant to think aloud while interacting with the DIARYPLAY authoring interface. With participant consent, we recorded the on-screen actions and took observation notes of their authoring behaviors. We

scheduled up to 30 minutes for the authoring task and informed participants that they could work at their own pace within this time frame. All participants completed the task within 20 minutes.

After the participant completed the authoring task, we conducted a semi-structured interview (around 20 minutes) about their authoring experience. We asked about their overall experience, perceived workload, the effectiveness of system features in authoring assistance, satisfaction with the created interactive vignette, and suggestions for system improvements in supporting interactive vignette creation for everyday storytelling.

Then, we invited the two participants to experience each other's interactive vignette as players. We informed participants that they need to play the interactive vignette from beginning to end, but were free to replay it as many times as they wished. We then ended the viewing task to ensure that the final half hour remained available for the remaining tasks. We asked the player to report their takeaway messages of the characters, events, environment, and feelings after the viewing experience, and we asked the author to comment on how well it aligns with their intended narrative.

Finally, we conducted a semi-structured interview (around 20 minutes) about participants' viewing experience. We asked what they liked and disliked about the interactive vignette delivered by DIARYPLAY and invited suggestions for improving how the system presents interactive vignettes in everyday storytelling.

To analyze the data, we used open coding [36] to inductively analyze participants' behaviors and reflections, with particular attention to how they related to DGs but open to unexpected insights. The first, second, and third authors reviewed the entire dataset, which included (1) interview transcripts, (2) observational notes, and (3) screen recordings. Each author then independently coded the data using both *in vivo* and *pattern coding* approaches. After the initial coding, we organized the codes into preliminary themes by merging our interpretations. Through ongoing team discussions, we refined the themes, aligning them with our research questions.

6.3 Results

Overall, our findings show that DIARYPLAY enables authors to create interactive vignettes that express their intended everyday stories, while balancing an engaging viewing experience with manageable authoring effort. As illustrated in Figure 9, participants used DIARYPLAY to express a wide variety of stories, spanning across various locations (e.g., office, library, museum, home, theater, shopping mall), characters (e.g., family members, friends, colleagues, strangers, classmates), and types of events (e.g., work day, school morning routine, museum visit). This demonstrates the DIARYPLAY's capability to support diverse expressions of everyday stories.

6.3.1 Answering RQ1. Overall, participants perceived the duration and workload of creating an interactive vignette as appropriate for everyday storytelling, indicating that the system fulfilled DG1. Specifically, participants completed interactive vignettes in 12–19 minutes ($\mu = 16$) with DIARYPLAY, a duration they regarded as reasonable for everyday storytelling. They appreciated the system's automatic and mostly accurate transformation from text story to interactive vignette elements: it eased the challenge of starting from scratch and provided “a great foundation for further authoring” (P1).

Today, I went to the office. I had a meeting with my co-workers, and Minsu who is one of my co-worker was presenting about his idea. After the meeting, we ate coffee and snack in the resting room.



Today I went to the museum with my mother and father. I bought coffee for my father. My father went to toilet. My mother waited for me and my father. Finally, we met at the center of the museum and went to see the contents.



In the morning, I woke up early to prepare to go to school, but my roommates told me that the class was cancelled. Instead, we decided to make some popcorn and watch a movie together. After the movie, we decided to go back to sleep.



Today I went to the shopping mall located in the city center with my best friend. We met my friend's mother there. She bought us ice-cream and helped us find a cute dress.



I met my friend at the school cafeteria to have lunch. We both had Chinese food for lunch. As I arrived earlier, I was already starting my meal when she was ordering some food. We had great lunch, and left the cafeteria for the next class.



Yesterday I went to the theater with my gf. She said she wanted to have some coke, so I went to buy it. While carrying a cup of coke, I spilled it on the floor, which made me buy another cup.



Figure 9: Example interactive vignettes created by participants from the user study. Minor grammatical errors in the original stories have been revised for illustration.

However, the system's parsing of input stories did not always align with participants' intentions. Two participants encountered errors in how the system parsed events, stemming from ambiguity in the story inputs. P13's story input had "we ate chicken together," which the system parsed as two people having a meal, though he had intended three. P9's phrases "had a meeting" and "talked about research topics" were meant as two separate events, but were initially parsed as one. Despite such errors, both participants felt that corrections were easy and fast to make in the authoring interface. Notably, P13 initially attempted to resolve the issue by editing the input text itself—rewrote the story as "three of us ate chicken together" and expected the event element to update automatically while preserving existing content. This behavior points to the potential of a more fluid authoring workflow, where modifications to the input story seamlessly propagate into the interactive vignette.

Beyond error correction in events, DIARYPLAY also supported the participants to refine the generated environment and characters to

better express their intentions. In environment design, participants repositioned three to five objects to improve realism or visual appeal. For instance, P11 moved a sofa closer to the TV and aligned two chairs to make the space feel more "realistic and visually appealing". P14 noted: "I felt I was doing something creative while arranging the environment, but it wasn't burdensome. I didn't need to draw assets—just by the text input, moving, adding, or deleting objects, I saw my imagined space appear in reality." The existing modification features satisfied most participants, while a few participants (3/16) desired finer control over objects, such as rotation and resizing.

We observed that a substantial portion of authoring time, about half on average, was devoted to character building. P10 explained: "Text input alone is too weak to express characters, especially their mood and personality. Building the persona gave me a chance to develop them more deeply." Additionally, participants found the process enjoyable, with P11 describing it as the "most fun and creative step." When filling in persona attributes, many participants (9/16)

found it difficult to articulate personality traits without guidance. Thanks to the DIARYPLAY’s conversation simulations and trait suggestions, they felt it was getting easier to build the persona. Other than being directly used by the system as examples when generating NPC dialogues with the player, the conversation simulations helped authors obtain a clearer understanding of how the character was understood by DIARYPLAY, and they can make further modifications to the persona. Moreover, the persona trait suggestions based on the conversation simulations helped authors fill in the persona form. P4, for example, initially struggled to “pick the right words to describe a friend’s personality beyond bland ‘good’ or ‘bad’”, but through conversation simulations, the system suggested the trait “supportive,” which he found to be highly relevant and used it. Beyond the current system design, some participants (4/16) expressed a desire for more nuanced emotional expression. Currently, authors can assign only a static “current mood” to a character, but participants wanted support for modeling dynamic emotional progression. As P5 put it: *“I wish I could design how a character’s emotions change in the story. For example, happy at first, then sad, then happy again.”* and suggested a feature for dynamic emotional modeling that evolves in sync with story events.

6.3.2 Answering RQ2. Through reviewing their own interactive vignettes and hearing takeaways from paired players, participants felt that DIARYPLAY successfully captured and conveyed their core story message by reflecting character personas and the intended sequence of events within the designed environment, thereby satisfying DG2. After completing the authoring process, participants previewed their interactive vignettes and confirmed that the outcome preserved their intended message, accurately reflecting key events and character persona. For example, P14 remarked, *“I’m satisfied that my character speaks the way I defined. I specified in the persona that she often says ‘oh really,’ and now her dialogue reflects that perfectly.”* Participants also felt their intentions were preserved when players shared their takeaways, even though each player engaged in at least one divergent activity. For instance, P7 viewed P8’s interactive vignette about three schoolmates meeting for a project in the library. Instead of joining the meeting immediately, P7 had the main character read a book from a shelf, while observing the two NPCs “turning on the monitor” and “preparing slides” in the meeting room. He then joined the meeting and talked with the NPCs, still following the core storyline. As P8 reflected, *“player understood what happened in my story, though there were divergent interactions. I think they are tiny but nice decorations. And the player also learned well about the characters, like their speaking tones and personality, through chatting with the NPCs.”*

At the same time, authors expressed a desire to review a play log of how players interacted with their interactive vignettes. P8 noted: *“I’m quite curious what exactly happened on the player’s end. I felt that as the author, I should be aware of what my audience experienced, even if the core events remained the same.”* Envisioning real-world use where their interactive vignettes might be viewed by a broader audience, P15 similarly expressed a desire to collect and review all viewing experiences, describing this as an *“author–player interaction, similar to social media posts.”*

Participants also found the original text story helpful for understanding the author-defined main storyline. However, a few

participants (P3, 5, 10) noted mismatches between the text story and what was actually unfolding in the interactive vignette. This discrepancy arose because the text story included the key events, omitting the divergent paths introduced during interaction. As P10 remarked, *“I felt the text story was somewhat of an overview, but not a real caption.”* At the same time, participants expressed the need to see a transparent distinction between human-authored and system-generated content. As P10 noted, *“Though NPC’s activities and their dialogues were guided by the human author’s intent, they were system-generated, not human-generated directly. As a player, I would like to clearly know that. It should note which parts are system-generated and which are human-made, to avoid misunderstanding.”* While participants generally appreciated the human-AI co-creative nature of the interactive vignettes, they felt that a lack of transparent authorship designation could lead to confusion.

6.3.3 Answering RQ3. All participants described interactive vignettes as engaging and valued them as a playful, memorable medium for everyday storytelling. They especially enjoyed the liveliness of NPCs and the chance to explore alternative branches within the same story, demonstrating that the system successfully met DG3. Participants appreciated the novelty of stepping into the author’s perspective and actively exploring the everyday experience through unfolding events. Most participants (14/16) described this format as more immersive and emotionally resonant than typical social media content such as text, images, or videos. As P15 explained, *“On social media, I usually just browse others’ daily updates quickly, but viewing the interactive vignette, I felt I was actually experiencing their daily moments, which let me resonate with it a lot.”*

Participants especially appreciated the liveliness of NPCs, who remained active and responsive even during player-initiated divergent activities. P11 noted, *“In RPGs I’ve played, some NPCs just stand there and feel oddly lifeless, but in my viewing experience, I felt like all the characters were thinking and alive.”* Many participants (12/16) revisited the interactive vignette multiple times to explore alternative interactions between key events. As P7 noted: *“It’s fun to explore these little deviations, like opening up a few multiverses while still sharing the same core storyline.”* Similarly, P14 tried different conversations across playthroughs and described the experience as *“feeling like living inside the story.”*

7 Discussion

We discuss (1) the balance of human authoring and system generation in everyday storytelling, (2) opportunities and considerations for integrating DiaryPlay into social media, and (3) ethical considerations for AI-generated content in everyday storytelling.

7.1 Balancing Human Authoring and System Generation in Everyday Storytelling

From the design and evaluation of DIARYPLAY, we derive several key design lessons. First, everyday storytellers should be able to modify intermediate results generated by the system. Our user study findings confirmed that the everyday storytellers found it useful to see and modify the system-generated narrative elements as intermediate results. Initially, we designed the interactive vignette specification dashboard as a way to surface the information gap between the text story input and what is required to construct

an interactive vignette (e.g., missing character personas or unspecified event objects). In practice, user study participants used the specification not only to fill in missing information, but also to verify whether the system’s interpretation aligned with their intent and to correct misrepresentations. This aligns with prior findings on human–AI cocreation [82], which highlight the importance of exposing AI-generated intermediate results to support iterative refinement and to enhance human authors’ visibility, flexibility, and sense of agency.

At the same time, everyday storytellers should be able to direct the authoring of interactive vignettes without needing to specify every narrative detail. User study participants appreciated being able to define key events and characters, and let the system automatically generate adaptive NPC activities and dialogues at interactive vignette runtime. This authoring approach enabled everyday storytellers to support player divergence within an intended narrative boundary (DG2, DG3), without requiring them to manually craft every narrative branch (DG1). Traditionally, supporting divergent player interactions has required extensive multi-branch authoring [10, 47, 50, 53], which emphasizes authors’ full control over every narrative branch and is incompatible with lightweight everyday storytelling. In contrast, the CD Module enabled a new human–AI co-authoring approach that balances author agency with low authoring effort: authors define the narrative boundary, and the system generates adaptive branches within it. This approach can extend beyond NPC behavior generation in interactive vignettes to other forms of everyday storytelling, such as generating text-based branches for interactive text stories [40] or image-based story material for interactive collages [9].

Importantly, the balance between the human authoring and system generation—as demonstrated in DIARYPLAY—was intentionally tailored for the lightweight, informal nature of everyday storytelling. The CD Module worked well because everyday storytellers were generally open to multiple narrative divergences as long as they remained consistent with the overall storyline logic and the defined character personas. For other storytelling domains, the balance between human authoring and system generation must be adjusted to fit the target context. For example, in interactive narrative scripting for product-level games, authors are more devoted users and often require strict control over narrative logic (e.g., ensuring an event occurs only after a specific trigger [57, 58]). In such contexts, while the human–AI coauthoring enabled by the CD Module may still be useful for automatically generating narrative branch options for inspiration and ideation, systems must offer more explicit control over branching structures. Authors should be able to carefully design and test each narrative path or define strict divergence conditions, rather than relying primarily on automated branching generation as in the current CD Module.

7.2 Opportunities and Considerations for Integrating DIARYPLAY into Social Media

To apply DIARYPLAY on real-world social media platforms, we discuss opportunities and considerations related to its interplay with conventional media, adaptation to mobile devices, and potential social interactions within the user community.

7.2.1 Combining Interactive Vignettes with Conventional Media for Everyday Storytelling. Building on the complementary expressive strengths of interactive vignettes and conventional media, we envision two ways of integrating these formats in everyday storytelling on social media.

First, interactive vignettes can serve as interactive annotations for text or video content. For example, social media users can use interactive vignettes to highlight key moments in blogs or vlogs (P12) and invite the audience to role-play, thereby adding playfulness and engagement to conventional everyday storytelling. Second, text, images, and videos can be incorporated into interactive vignettes. To better convey narrative progression, the interactive vignette could employ text captions that dynamically highlight words or phrases associated with the current event, helping ground players’ understanding of unfolding actions [37, 38]. Beyond using text as captions, social media users could embed personal photos or video clips as storyworld assets (e.g., a family photo on the wall, a video playing on a TV), enhancing realism and grounding the interactive vignette in everyday life.

7.2.2 Adaptation to Mobile Devices. User study participants emphasized that, in everyday use, they would primarily engage with interactive vignettes on mobile devices [29], potentially through direct sharing or through posting on existing social media platforms (e.g., X [101], Instagram [31]). The current version of DIARYPLAY is designed for a desktop environment, but mobile devices with smaller screens require additional design considerations for both authors and players.

For authors, the limited screen size of mobile devices cannot accommodate the three panels (environment, characters, and events) simultaneously. A potential design adaptation is a sequential workflow that guides authors through environment, character, and event creation step by step, with each panel displayed on a single screen and the option to revisit earlier steps. For players, small screens may make details harder to see. Possible adaptations include zooming functions or a following-camera perspective on the main character to support fluid navigation within constrained display spaces. These considerations suggest that mobile-first design will be essential for deploying DIARYPLAY in real-world social media contexts.

7.2.3 Social Interactions within the User Community. Social media is not only about publishing and viewing content but also about fostering author–player social interactions [5]. User study participants (P8, 15) expressed curiosity about players’ experiences and reactions, suggesting features such as replay logs or comment mechanisms to support two-way communication. For instance, the author can review players’ playthroughs and reply to players’ comments to better communicate how the everyday stories were experienced from players’ perspectives. Additionally, to enable exploration from different perspectives, interactive vignettes can allow players to choose from multiple characters to role-play, or allow multiple players to co-participate in the same vignette.

Among authors, we also see potential for a user community, which shares user-generated, open-sourced assets and templates for interactive vignettes, similar to user-generated Instagram AR filters [30]. This user community could expand the diversity of visual styles, characters, and environments available for authoring.

Incorporating such collaborative practices could expand the expressive range of interactive vignettes while aligning with familiar social media conventions.

7.3 Ethical Considerations for AI-Generated Content in Everyday Storytelling

Despite DIARYPLAY introducing the exciting opportunity to adopt interactive vignettes in everyday storytelling, we must also consider potential ethical issues that should be recognized and carefully addressed through design.

7.3.1 Risk of Harmful or Biased Content. Although OpenAI’s usage policy helped to filter some harmful content, interactive vignettes could still be misused if authored or viewed with malicious intent. Characters modeled after real people may be deliberately scripted or manipulated to behave harmfully—for example, by speaking offensively, spreading misinformation, or engaging in actions that damage reputations. In addition, the use of LLMs introduces the risk of biased or stereotypical outputs, especially in generated conversations [99]. Such biases could reinforce harmful narratives or marginalize certain groups if left unchecked. We believe that deployment on social media would leverage the established moderation approaches [91]—including flagging, review workflows, and appeals—to ensure responsible use and safeguard users from harmful or biased content.

7.3.2 Risk of Implanting False Memory. Although participants in both the formative and user studies viewed interactive vignettes not as 1:1 replicas of daily life but as interactive narratives with divergent events and dialogues, DIARYPLAY might risk implanting false memories [67] if players mistake AI-generated content as real interactions. To mitigate this, we emphasize system transparency around human–AI authorship. The system should clearly indicate where AI-generated content is involved and prompt players to engage critically. This approach is similar to social media practices for deepfakes [21], where posts are labeled with warnings such as “this post contains AI-generated content, please discern carefully.”

8 Limitations and Future Work

Limitations of the User Study. While DIARYPLAY gathered in-depth qualitative feedback on the authoring and viewing experience, the user study has several limitations. First, it was limited to one-time interactions, and a longitudinal deployment (e.g., over several weeks) would provide insights into how DIARYPLAY supports everyday storytelling over time (e.g., how familiarity with the system influences their authoring experience). Second, our participant pool was culturally homogeneous, leaving future work to explore how storytelling practices and system use may vary across cultural contexts [56]. Third, the study took place in a controlled lab setting, where the presence of researchers may have influenced participants’ behaviors. Future work should therefore include in-the-wild, cross-cultural, and longer-term studies to better understand how DIARYPLAY integrates into everyday storytelling practices. To enable long-term, repeated use of DIARYPLAY, the system would need features that let users reuse environment templates and characters.

Extending DIARYPLAY’s Story Scope. The current version of DIARYPLAY focuses on simple, event-centric everyday stories; future

extensions could broaden its story scope to include more characters, transitions across multiple locations, and richer emotional expression. Supporting stories with more characters and transitions between spatially discontinuous locations (e.g., moving from home to school) would allow DIARYPLAY to represent more complicated and richer slices of everyday life. Enhancing the system’s ability to capture emotional expression would further expand its expressive range, which requires extending the interactive vignette specification to include characters’ emotional description, their evolution across the events, and their influence on character behaviors over time. However, these extensions increase authoring complexity, as users must manage a larger set of narrative elements. Future systems should therefore explore intelligent interface support, such as LLM-assisted timeline-based authoring tools [58], to keep authoring effort manageable while still enabling flexible editing across multiple story components.

Extending DIARYPLAY’ Viewing Experience. While participants appreciated the current viewing experience of interactive vignettes in DIARYPLAY, future work can extend its social aspects and spatial expressiveness: allowing players to choose different characters, supporting multiplayer co-participation, and expanding from 2D to 3D spatial storytelling. Supporting character selection would require authors to prepare core narrative paths for each character, ensuring that switching characters provides meaningfully different role-play perspectives. Enabling multiplayer co-participation introduces a different challenge: the system will need to coordinate multiple players while still maintaining the author’s intended storyline, rather than drifting into a fully open-ended narrative. At the same time, expanding to 3D storytelling poses significant challenges. Although recent text-to-3D advances (e.g., Genie 3 [66], SceneTeller [61], LayoutGPT [18]) demonstrate the feasibility of generating complex 3D worlds from narrative text, systems must still ensure that generated scenes include all story-relevant objects and relations, and remain easily editable so authors can refine layouts and maintain creative control. A 3D environment may also raise expectations for realistic character behaviors (e.g., cooking, eating, jumping), further increasing authoring effort. Realizing these extensions will therefore require future systems to provide effective generation supports and lightweight refinement tools to avoid overwhelming everyday storytellers.

Potential Applications Beyond Everyday Storytelling. While DIARYPLAY was designed to support everyday storytelling, its low-effort authoring process, roleplay mechanics, and the branch-and-bottleneck narrative structure suggest potential for broader applications. First, in the interactive media industry [16], DIARYPLAY could serve as a rapid prototyping tool for designers to test and communicate ideas visually. Second, in mental health counseling [79], interactive vignettes might be customized to support ad hoc interactive narratives, though such use would require careful ethical and clinical consideration. Third, in education [83], adapting the system for age-appropriate, curriculum-aligned content could position interactive vignettes as an engaging learning tool. Extending DIARYPLAY beyond everyday storytelling would require rethinking the vignette specification, integrating domain-specific constraints, and aligning with professional practices. Collaboration with domain experts, such as game designers, clinicians, and educators,

will be essential to guide appropriate adaptation and evaluation. These opportunities point toward promising directions for future research and development.

9 Conclusion

In this work, we introduced DIARYPLAY, an AI-assisted authoring system that enables lightweight creation of interactive vignettes for everyday storytelling. Interactive vignettes allow the audience to role-play the author's perspective and actively participate in the story, which motivated social media users to adopt the format into everyday storytelling. However, authoring interactive vignettes using existing authoring tools is technically demanding and time-intensive, which conflicts with the immediacy required in everyday storytelling. At the same time, the aspiration to provide branching narratives with adaptive NPC behaviors that allow divergent activities beyond the core storyline further increases authoring complexity. DIARYPLAY addresses these challenges by automatically transforming text stories into interactive vignette elements and guiding authors in refinement, and using the CD Module to automatically plan NPC behaviors that adapt to player interactions while preserving the author's intended story message. A preliminary technical evaluation showed that the module generates believable NPC activities based on the author-defined character personas and the overarching storyline. A user study demonstrated that DIARYPLAY supports low-effort authoring while enabling engaging viewing experiences that convey the core story message. Finally, we discuss the real-world applicability of DIARYPLAY on social media platforms and highlight potential extensions beyond everyday storytelling. Overall, while our research is among the early efforts toward enabling interactive vignette creation in everyday use with several limitations, we believe that the system hints at a promising new form of social everyday storytelling and offers insights and agendas for future work.

Acknowledgments

This work was supported by a grant of the KAIST-KT joint research project funded by KT (G01220645), the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No. 2021-0-01347, Video Interaction Technologies Using Object-Oriented Video Modeling), the Culture, Sports and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture, Sports and Tourism in 2026 (Project Name: Development of Multimodal UX Evaluation Platform Fechnology for XR Spatial Responsive Content Optimization, Grant No.: RS-2024-00361757, Contribution Rate: 20%), and the 2025-1 Yonsei University Future-Leading Research Initiative (Grant No.: 2025-22-0156). We used ChatGPT [64] to perform grammar checks, as well as obtain phrase suggestions and grammatical restructuring based on our own texts.

References

- [1] Awal Ahmed Fime, Saifuddin Mahmud, Arpita Das, Md. Sunzidul Islam, and Jong-Hoon Kim. 2025. Automatic Scene Generation: State-Of-The-Art Techniques, Models, Datasets, Challenges, and Future Prospects. *IEEE Access* 13 (2025), 95753–95796. <https://doi.org/10.1109/ACCESS.2025.3574298>
- [2] Rehaf Aljammaz, Noah Wardrip-Fruin, and Michael Mateas. 2023. Towards an Understanding of Character Believability. In *Proceedings of the 18th International Conference on the Foundations of Digital Games*. Association for Computing Machinery, New York, NY, USA, Article 24, 9 pages. <https://doi.org/10.1145/3582437.3582466>
- [3] Ruth Aylett. 2000. Emergent Narrative, Social Immersion and “Storification”. In *Proceedings of the 1st international workshop on narrative and interactive learning environments*, Vol. 8. Sense Publishers, Rotterdam, The Netherlands, 35–44.
- [4] Ruth Aylett, Sandy Louchart, and Allan Weallans. 2011. Research in Interactive Drama Environments, Role-Play and Story-Telling. In *Interactive Storytelling*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1–12. https://doi.org/10.1007/978-3-642-25289-1_1
- [5] Ava Bartolome and Shuo Niu. 2023. A Literature Review of Video-Sharing Platform Research in HCI. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 790, 20 pages. <https://doi.org/10.1145/3544548.3581107>
- [6] Eden Bensaid, Mauro Martino, Benjamin Hoover, and Hendrik Strobelt. 2021. FairyTailor: A Multimodal Generative Framework for Storytelling. <https://doi.org/10.48550/arXiv.2108.04324>
- [7] Anton Bogdanovych, Tomas Trescak, and Simeon Simoff. 2016. What Makes Virtual Agents Believable? *Connection Science* 28, 1 (2016), 83–108. <https://doi.org/10.1080/09540091.2015.1130021>
- [8] Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, Yusuf Aytar, Sarah Maria Elisabeth Bechtle, Feryal Behbahani, Stephanie C.Y. Chan, Nicolas Heess, Lucy Gonzalez, Simon Osindero, Sherjil Ozair, Scott Reed, Jingwei Zhang, Konrad Zolna, Jeff Clune, Nando De Freitas, Satinder Singh, and Tim Rocktäschel. 2024. Genie: Generative Interactive Environments. In *Proceedings of the 41st International Conference on Machine Learning*, Vol. 235. Microtome Publishing, Brookline, MA, Article 185, 21 pages. <https://proceedings.mlr.press/v235/bruce24a.html>
- [9] V. Buckenham. 2022. Downpour. Retrieved November 18, 2025 from <https://downpour.games/>
- [10] Alex Calderwood, Noah Wardrip-Fruin, and Michael Mateas. 2022. Spinning Coherent Interactive Fiction Through Foundation Model Prompts. In *Proceedings of the 13th International Conference on Computational Creativity*. Association for Computational Creativity, Bozen-Bolzano, Italy, 44–53. http://computationalcreativity.net/iccc22/papers/ICCC-2022_paper_65.pdf
- [11] Yee Bee Choo, Tina Abdullah, and Abdullah Mohd Nawil. 2020. Digital Storytelling vs. Oral Storytelling: An Analysis of the Art of Telling Stories Now and Then. *Universal Journal of Educational Research* 8, 5A (2020), 46–50. <https://doi.org/10.13189/ujer.2020.081907>
- [12] John Joon Young Chung and Max Kreminski. 2024. Patchview: LLM-Powered Worldbuilding With Generative Dust and Magnet Visualization. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, Article 77, 19 pages. <https://doi.org/10.1145/3654777.3676352>
- [13] Kate Compton and Michael Mateas. 2015. Casual Creators. In *Proceedings of the Sixth International Conference on Computational Creativity (ICCC 2015)*. Brigham Young University, Park City, Utah, 228–235. http://computationalcreativity.net/iccc2015/proceedings/10_2Compton.pdf
- [14] Cassidy Curtis, Sigurdur Orn Adalgeirsson, Horia Stefan Ciurdar, Peter McDermott, JD Velásquez, W. Bradley Knox, Alonso Martinez, Dei Gaztelumendi, Norberto Adrian Goussies, Tianyu Liu, and Palash Nandy. 2022. Toward Believable Acting for Autonomous Animated Characters. In *Proceedings of the 15th ACM SIGGRAPH Conference on Motion, Interaction and Games*. Association for Computing Machinery, New York, NY, USA, Article 1, 15 pages. <https://doi.org/10.1145/3561975.3562941>
- [15] Adam Le Doux. 2017. Bitsy Game Maker. Retrieved November 18, 2025 from <https://www.bitsy.org/>
- [16] Henrik Engström, Jenny Brusik, and Patrik Erlandsson. 2018. Prototyping tools for game writers. *The Computer Games Journal* 7, 3 (2018), 153–172.
- [17] Min Fan, Xinyue Cui, Wanqing Ma, Haiyan Li, Xin Tong, Lin Yang, and Yonghui Wang. 2025. From Words to Wonder: Designing and Evaluating an AI-Empowered Creative Storytelling System for Elementary Children. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 328, 15 pages. <https://doi.org/10.1145/3706598.3713478>
- [18] Weixi Feng, Wanrong Zhu, Tsu-jui Fu, Varun Jampani, Arjun Akula, Xuehai He, Sugato Basu, Xin Eric Wang, and William Yang Wang. 2023. Layout-GPT: Compositional Visual Planning and Generation With Large Language Models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, Article 802, 26 pages. https://proceedings.neurips.cc/paper_files/paper/2023/file/3a7f9e485845dac27423375c934cb4db-Paper-Conference.pdf
- [19] Joshua A. Fisher. 2023. Centering the Human: Digital Humanism and the Practice of Using Generative AI in the Authoring of Interactive Digital Narratives. In *Interactive Storytelling*. Springer Nature Switzerland, Cham, 73–88. https://doi.org/10.1007/978-3-031-47655-6_5
- [20] FromSoftware. 2022. Elden Ring. Retrieved November 25, 2025 from <https://en.bandainamcoent.eu/elden-ring/elden-ring>

- [21] Dilrukshi Gamage, Dilki Sewwandi, Min Zhang, and Arosha K Bandara. 2025. Labeling Synthetic Content: User Perceptions of Label Designs for AI-Generated Content on Social Media. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 814, 29 pages. <https://doi.org/10.1145/3706598.3713171>
- [22] Game Freak. 1996. Pokémon Red and Pokémon Blue. Video game.
- [23] Paulo Gomes, Ana Paiva, Carlos Martinho, and Arnab Jhala. 2013. Metrics for Character Believability in Interactive Narrative. In *Interactive Storytelling*. Springer International Publishing, Cham, 223–228. https://doi.org/10.1007/978-3-319-02756-2_27
- [24] Jack Goody. 2006. From Oral to Written: An Anthropological Breakthrough in Storytelling. *The novel* 1 (2006), 3–36. <https://doi.org/10.1515/9780691243757-002>
- [25] Eric Gordon and Steven Schirra. 2011. Playing With Empathy: Digital Role-Playing Games in Public Meetings. In *Proceedings of the 5th International Conference on Communities and Technologies*. Association for Computing Machinery, New York, NY, USA, 179–185. <https://doi.org/10.1145/2103354.2103378>
- [26] Christine Greenhow and Trevor Owens. 2011. Social Videogame Creation: Lessons From RPG Maker. *On the Horizon: The International Journal of Learning Futures* 19, 1 (02 2011), 52–61. <https://doi.org/10.1108/10748121111107708>
- [27] Stefan Göbel, Luca Salvatore, and Robert Konrad. 2008. StoryTec: A Digital Storytelling Platform for the Authoring and Experiencing of Interactive and Non-Linear Stories. In *2008 International Conference on Automated Solutions for Cross Media Content and Multi-Channel Distribution*. IEEE Computer Society, Washington, DC, USA, 103–110. <https://doi.org/10.1109/AXMEDIS.2008.45>
- [28] Chris Hales. 2015. Interactive Cinema in the Digital Age. In *Interactive Digital Narrative*. Routledge, New York, NY, USA, 36–50. <https://doi.org/10.4324/9781315769189>
- [29] Lee Humphreys. 2013. Mobile Social Media: Future Challenges and Opportunities. *Mobile Media & Communication* 1, 1 (2013), 20–25. <https://doi.org/10.1177/2050157912459499>
- [30] Sergio Ibáñez-Sánchez, Carlos Orus, and Carlos Flavian. 2022. Augmented Reality Filters on Social Media. Analyzing the Drivers of Playability Based on Uses and Gratifications Theory. *Psychology & Marketing* 39, 3 (2022), 559–578. <https://doi.org/10.1002/mar.21639>
- [31] Instagram, Inc. 2010. Instagram. Retrieved September 9, 2024 from <https://www.instagram.com>
- [32] itch.io. 2013. Retrieved September 8, 2025 from <https://itch.io/>
- [33] Kadokawa Corporation. 2024. RPG Maker. Retrieved May 18, 2024 from <https://www.rpgmakerweb.com/>
- [34] Eric R. Kandel, John D. Koester, Sarah H. Mack, and Steven A. Siegelbaum. 2021. *Visual Processing for Attention and Action*. McGraw Hill, New York, NY, Chapter 25, 582–597. <https://neurology.mhmedical.com/content.aspx?aid=1180642180>
- [35] Aisling Kellihier and Glorianna Davenport. 2007. Everyday Storytelling: Supporting the Mediated Expression of Online Personal Testimony. In *Human-Computer Interaction. HCI Applications and Services*. Springer Berlin Heidelberg, Berlin, Heidelberg, 926–933. https://doi.org/10.1007/978-3-540-73111-5_102
- [36] Shahedul Huq Khandkar. 2009. Open Coding. *University of Calgary* 23, 2009 (2009), 2009.
- [37] Dae Hyun Kim, Enamul Hoque, Juho Kim, and Maneesh Agrawala. 2018. Facilitating Document Reading by Linking Text and Tables. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 423–434. <https://doi.org/10.1145/3242587.3242617>
- [38] Dae Hyun Kim, Vidya Setlur, and Maneesh Agrawala. 2021. Towards Understanding How Readers Integrate Charts and Captions: A Case Study with Line Charts. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 610, 11 pages. <https://doi.org/10.1145/3411764.3445443>
- [39] Taewan Kim, Donghoon Shin, Young-Ho Kim, and Hwajung Hong. 2024. Diary-Mate: Understanding User Perceptions and Experience in Human-AI Collaboration for Personal Journaling. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 1046, 15 pages. <https://doi.org/10.1145/3613904.3642693>
- [40] Chris Klimas. 2009. *Twine*. <https://twinery.org/> Version 2.3.14.
- [41] Hartmut Koenitz and Sandy Louchart. 2015. Practicalities and Ideologies, (Re)-Considering the Interactive Digital Narrative Authoring Paradigm. In *Proceedings of the 10th International Conference on the Foundations of Digital Games*. Association for Computing Machinery, New York, NY, USA, 1–4. http://www.fdg2015.org/papers/fdg2015_paper_03.pdf
- [42] Anastasia Konstantopoulou, Eleni Nikolou, Georgios Fesakis, Stamatia Volika, and Georgios Markogiannakis. 2018. Designing Interactive Digital Storytelling as a Strategy of Raising Children’s Awareness of Bullying in Preschool Education: Implications for Bullying Prevention. In *Proceedings of the International Digital Storytelling Conference DST 2018*. 91. <https://dst.ntlab.gr/2018/proceedings/>
- [43] Max Kreminski, Melanie Dickinson, Joseph Osborn, Adam Summerville, Michael Mateas, and Noah Wardrip-Fruin. 2020. Germinate: A Mixed-Initiative Casual Creator for Rhetorical Games. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 16, 1 (Oct. 2020), 102–108. <https://doi.org/10.1609/aiide.v16i1.7417>
- [44] Vikram Kumaran, Jonathan Rowe, and James Lester. 2024. NarrativeGenie: Generating Narrative Beats and Dynamic Storytelling With Large Language Models. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 20 (Nov. 2024), 76–86. <https://doi.org/10.1609/aiide.v20i1.31868>
- [45] Vikram Kumaran, Jonathan Rowe, Bradford Mott, and James Lester. 2023. SceneCRAFT: Automating Interactive Narrative Scene Generation in Digital Games With Large Language Models. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*. AAAI Press, Palo Alto, CA, 86–96. <https://doi.org/10.1609/aiide.v19i1.27504>
- [46] William Labov and Joshua Waletzky. 1997. Narrative Analysis: Oral Versions of Personal Experience. *Journal of Narrative & Life History* 7, 1–4 (1997), 3–38. <https://doi.org/10.1075/jnlh.7.02nar>
- [47] Jorge Leandro, Sudha Rao, Michael Xu, Weijia Xu, Nebojsa Jovic, Chris Brockett, and Bill Dolan. 2024. GENEVA: GENERating and Visualizing Branching Narratives Using LLMs. In *2024 IEEE Conference on Games (CoG)*. IEEE, Piscataway, NJ, USA, 1–5. <https://doi.org/10.1109/CoG60054.2024.10645625>
- [48] Michael Sangyeob Lee and Carrie Heeter. 2012. What Do You Mean by Believable Characters?: The Effect of Character Rating and Hostility on the Perception of Character Believability. *Journal of Gaming & Virtual Worlds* 4, 1 (2012), 81–97. https://doi.org/10.1386/jgvw.4.1.81_1
- [49] LibreGameWiki. 2022. Sleep Is Death. Retrieved November 18, 2025 from https://libregamewiki.org/Sleep_Is_Death
- [50] M. T. Linaza, H. Eskudero, C. Lamsfus, and G. Marcos. 2004. An Authoring Tool for Interactive Digital Storytelling. In *Proceedings of the 5th International Conference on Virtual Reality, Archaeology and Intelligent Cultural Heritage*. Eurographics Association, Goslar, DEU, 203–211. <https://doi.org/10.2312/VAST/VAST04/203-211>
- [51] Ryan Louie, Jesse Engel, and Cheng-Zhi Anna Huang. 2022. Expressive Communication: Evaluating Developments in Generative Models and Steering Interfaces for Music Creation. In *Proceedings of the 27th International Conference on Intelligent User Interfaces*. Association for Computing Machinery, New York, NY, USA, 405–417. <https://doi.org/10.1145/3490099.3511159>
- [52] A Bryan Loyall. 1997. *Believable Agents: Building Interactive Personalities*. Ph.D. Dissertation. Carnegie Mellon University.
- [53] Zhuoran Lu, Qian Zhou, and Yi Wang. 2025. WhatELSE: Shaping Narrative Spaces at Configurable Level of Abstraction for AI-bridged Interactive Storytelling. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 333, 18 pages. <https://doi.org/10.1145/3706598.3713363>
- [54] Michael Mateas and Andrew Stern. 2003. Façade: An Experiment in Building a Fully-Realized Interactive Drama. In *Game developers conference*, Vol. 2. Citeseer, 4–8. <https://users.soe.ucsc.edu/~michaelm/publications/mateas-gdc2003.pdf>
- [55] Maxis. 2000. Bloodborne. Retrieved November 25, 2025 from <https://en.wikipedia.org/wiki/Bloodborne>
- [56] Dan P McAdams. 2019. “First We Invented Stories, Then They Changed Us”: The Evolution of Narrative Identity. *Evolutionary Studies in Imaginative Culture* 3, 1 (2019), 1–18. <https://doi.org/10.26613/esic.3.1.110>
- [57] Michael Mateas, Andrew Stern. 2005. Façade. Interactive Drama.
- [58] Aditi Mishra, Frederik Brudy, Qian Zhou, George Fitzmaurice, and Fraser Anderson. 2025. WhatIF: Branched Narrative Fiction Visualization for Authoring Emergent Narratives Using Large Language Models. In *Proceedings of the 2025 Conference on Creativity and Cognition*. Association for Computing Machinery, New York, NY, USA, 590–605. <https://doi.org/10.1145/3698061.3726933>
- [59] Mojang Studios. 2011. Minecraft. Retrieved September 11, 2024 from <https://www.minecraft.net/>
- [60] Noam Shazeer and Daniel de Freitas, co-founders. 2025. Character.AI: A Generative AI Chatbot Platform. Retrieved September 5, 2025 from <https://character.ai/>
- [61] Başak Melis Öcal, Maxim Tatarchenko, Sezer Karaoğlu, and Theo Gevers. 2024. SceneTeller: Language-to-3D Scene Generation. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LXXXV*. Springer-Verlag, Berlin, Heidelberg, 362–378. https://doi.org/10.1007/978-3-031-73013-9_21
- [62] OpenAI. 2022. DALL-E 2. Retrieved September 8, 2025 from <https://openai.com/research/dall-e-2>
- [63] OpenAI. 2024. GPT-4o. Retrieved September 8, 2025 from <https://platform.openai.com/docs/models/gpt-4o>
- [64] OpenAI. 2025. ChatGPT. Retrieved March 4, 2025 from <https://chatgpt.com/>
- [65] Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, Article 2, 22 pages. <https://doi.org/10.1145/3586183.3606763>
- [66] Jack Parker-Holder and Shlomi Fruchter. 2025. Genie 3: A New Frontier for World Models. Retrieved August 8, 2025 from <https://deepmind.google/discover/blog/genie-3-a-new-frontier-for-world-models/>

- [67] Pat Pataranutaporn, Chayapat Archiwangunprok, Samantha W. T. Chan, Elizabeth Loftus, and Pattie Maes. 2025. Synthetic Human Memories: AI-Edited Images and Videos Can Implant False Memories and Distort Recollection. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 538, 20 pages. <https://doi.org/10.1145/3706598.3713697>
- [68] Federico Peinado and Pablo Gervás. 2004. Transferring Game Mastering Laws to Interactive Digital Storytelling. In *Technologies for Interactive Digital Storytelling and Entertainment*. Springer Berlin Heidelberg, Berlin, Heidelberg, 48–54. https://doi.org/10.1007/978-3-540-27797-2_7
- [69] Rodrigo Perez-Vega, Kathryn Waite, and Kevin O’Gorman. 2016. Social Impact Theory: An Examination of How Immediacy Operates as an Influence Upon Social Media Interaction in Facebook Fan Pages. *The Marketing Review* 16, 3 (2016), 299–321. <https://doi.org/10.1362/146934716X14636478977791>
- [70] Photon Storm Ltd. 2013. Phaser: HTML5 Game Framework. <https://phaser.io/>. Accessed: 2025-09-08.
- [71] Ioannis A. Pissourios and Alexandros Ph. Lagopoulos. 2017. The Classification of Urban Uses. *Urban Science* 1, 3 (2017), 26 pages. <https://doi.org/10.3390/urbansci1030026>
- [72] Edward Powley, Mark Nelson, Swen Gaudl, Simon Colton, Blanca Pérez Ferrer, Rob Saunders, Peter Ivey, and Michael Cook. 2017. Wewva: Democratising Game Design. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 13, 1 (Jun. 2017), 273–275. <https://doi.org/10.1609/aiide.v13i1.12920>
- [73] PuzzleScript. 2013. Retrieved November 18, 2025 from <https://www.puzzlescript.net/>
- [74] Hua Xuan Qin, Shan Jin, Ze Gao, Mingming Fan, and Pan Hui. 2024. CharacterMeet: Supporting Creative Writers’ Entire Story Character Construction Processes Through Conversation With Llm-Powered Chatbot Avatars. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 1051, 19 pages. <https://doi.org/10.1145/3613904.3642105>
- [75] Anna Marie Rezk and Mads Haahr. 2022. Beyond Free Will: Understanding Approaches to Agency and Their Suitability for Bandersnatch-Like Titles. *Entertainment Computing* 43 (2022), 100500. <https://doi.org/10.1016/j.entcom.2022.100500>
- [76] Mark Owen Riedl and Vadim Bulitko. 2012. Interactive Narrative: An Intelligent Systems Approach. *AI Magazine* 34, 1 (Dec. 2012), 67. <https://doi.org/10.1609/aimag.v34i1.2449>
- [77] Mark O. Riedl and R. Michael Young. 2010. Narrative Planning: Balancing Plot and Character. *J. Artif. Int. Res.* 39, 1 (Sept. 2010), 217–268. <https://doi.org/10.1613/jair.2989>
- [78] Roblox. 2006. Retrieved November 24, 2025 from <https://www.roblox.com/>
- [79] Solrun Brenk Rønning and Stål Bjørkly. 2019. The Use of Clinical Role-Play and Reflection in Learning Therapeutic Communication Skills in Mental Health Education: An Integrative Review. *Advances in medical education and practice* 10 (2019), 415–425. <https://doi.org/10.2147/AMEP.S202115>
- [80] Rosebud AI. 2019. Rosebud AI: AI-Powered Game Development Platform. Retrieved September 12, 2024 from <https://www.rosebud.ai/>
- [81] Adrian L. Jessup Schneider and T.C. Nicholas Graham. 2017. Nudging and Shoving: Using In-Game Cues to Guide Player Exertion in Exergames. *Entertainment Computing* 19 (2017), 83–100. <https://doi.org/10.1016/j.entcom.2017.01.002>
- [82] Yang Shi, Tian Gao, Xiaohan Jiao, and Nan Cao. 2023. Understanding Design Collaboration Between Designers and Artificial Intelligence: A Systematic Literature Review. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2, Article 368 (Oct. 2023), 35 pages. <https://doi.org/10.1145/3610217>
- [83] Ulrike Spierling. 2006. Learning With Digital Agents—Integration of Simulations, Games, and Storytelling. In *Digital game based learning, Proceedings of the 4th international symposium for information design*. Universitätsverlag Karlsruhe, KIT Scientific Publishing, Karlsruhe, Germany, 115–148. <https://doi.org/10.5445/KSP/1000004726>
- [84] Ulrike Spierling and Nicolas Szilas. 2009. Authoring Issues Beyond Tools. In *Proceedings of the 2nd Joint International Conference on Interactive Digital Storytelling: Interactive Storytelling*. Springer-Verlag, Berlin, Heidelberg, 50–61. https://doi.org/10.1007/978-3-642-10643-9_9
- [85] Ulrike Martina Spierling. 2012. *Implicit Creation’–Non-Programmer Conceptual Models for Authoring in Interactive Digital Storytelling*. Ph. D. Dissertation. University of Plymouth.
- [86] Square. 1995. Chrono Trigger. Video game.
- [87] Jackie Strause. 2018. *Black Mirror: Bandersnatch - Netflix’s Interactive Film Explained*. The Hollywood Reporter. Retrieved November 28, 2025 from <https://www.hollywoodreporter.com/tv/tv-news/black-mirror-bandersnatch-netflixs-interactive-film-explained-1171486/>
- [88] Yuqian Sun, Xuran Ni, Haozhen Feng, Ray LC, Chang Hee Lee, and Ali Asadipour. 2022. Bringing Stories to Life in 1001 Nights: A Co-creative Text Adventure Game Using a Story Generation Model. In *Interactive Storytelling: 15th International Conference on Interactive Digital Storytelling, ICIDS 2022, Santa Cruz, CA, USA, December 4–7, 2022, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 651–672. https://doi.org/10.1007/978-3-031-22298-6_42
- [89] Frank Thomas. 1995. *The Illusion of Life*. Hyperion, New York, NY, USA.
- [90] Mike Treanor, Bryan Blackford, Michael Mateas, and Ian Bogost. 2012. Game-O-Matic: Generating Videogames That Represent Ideas. In *Proceedings of the The Third Workshop on Procedural Content Generation in Games*. Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/2538528.2538537>
- [91] Collins Udanor and Chinatu C. Anyanwu. 2019. Combating the Challenges of Social Media Hate Speech in a Polarized Society: A Twitter Ego Lexalytics Approach. *Data Technologies and Applications* 53, 4 (09 2019), 501–527. <https://doi.org/10.1108/DTA-01-2019-0007>
- [92] Unity. 2023. Retrieved September 8, 2025 from <https://unity.com/> Game development platform.
- [93] Jack Urbanek, Angela Fan, Siddharth Karamcheti, Saachi Jain, Samuel Humeau, Emily Dinan, Tim Rocktäschel, Douwe Kiela, Arthur Szlam, and Jason Weston. 2019. Learning to Speak and Act in a Fantasy Text Adventure Game. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 673–683. <https://doi.org/10.18653/v1/D19-1062>
- [94] Renato Verdugo, Miguel Nussbaum, Pablo Corro, Pablo Núñez Palma, and Paula Navarrete. 2011. Interactive Films and Construction. *ACM Trans. Multimedia Comput. Commun. Appl.* 7, 4, Article 39 (Dec. 2011), 24 pages. <https://doi.org/10.1145/2043612.2043617>
- [95] Elena Villaespesa and Sara Wolkowych. 2020. Ephemeral Storytelling With Social Media: Snapchat and Instagram Stories at the Brooklyn Museum. *Social Media + Society* 6, 1 (2020), 2056305119898776. <https://doi.org/10.1177/2056305119898776>
- [96] Yi Wang, Qian Zhou, and David Ledo. 2024. StoryVerse: Towards Co-Authoring Dynamic Plot With LLM-Based Character Simulation via Narrative Planning. In *Proceedings of the 19th International Conference on the Foundations of Digital Games*. Association for Computing Machinery, New York, NY, USA, Article 55, 4 pages. <https://doi.org/10.1145/3649921.3656987>
- [97] Stephen G. Ware and Cory Siler. 2021. Sabre: A Narrative Planner Supporting Intention and Deep Theory of Mind. In *Proceedings of the Seventeenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, Vol. 17. AAAI Press, Palo Alto, CA, Article 34, 8 pages. <https://doi.org/10.1609/aiide.v17i1.18896>
- [98] Stephen G. Ware and R. Michael Young. 2014. Glaiave: A State-Space Narrative Planner Supporting Intentionality and Conflict. In *Proceedings of the Tenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*. AAAI Press, Palo Alto, CA, 80–86. <https://doi.org/10.1609/aiide.v10i1.12712>
- [99] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. 2021. Ethical and Social Risks of Harm From Language Models. <https://doi.org/10.48550/arXiv.2112.04359>
- [100] Zhuohao Wu, Danwen Ji, Kaiwen Yu, Xianxu Zeng, Dingming Wu, and Mohamad Shidujaman. 2021. AI Creativity and the Human-AI Co-creation Model. In *Human-Computer Interaction. Theory, Methods and Tools*. Springer International Publishing, Cham, Switzerland, 171–190. https://doi.org/10.1007/978-3-030-78462-1_13
- [101] X (formerly Twitter). 2023. Retrieved September 8, 2025 from <https://x.com>
- [102] Yadong Xi, Xiaoxi Mao, Le Li, Lei Lin, Yanjiang Chen, Shuhan Yang, Xuhan Chen, Kailun Tao, Zhi Li, Gongzheng Li, Lin Jiang, Siyan Liu, Zeng Zhao, Minlie Huang, Changjie Fan, and Zhipeng Hu. 2021. KuiLeiXi: a Chinese Open-Ended Text Adventure Game. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 175–184. <https://doi.org/10.18653/v1/2021.acl-demo.21>
- [103] Yang Xiao, Yi Cheng, Jinlan Fu, Jiashuo Wang, Wenjie Li, and Pengfei Liu. 2023. How Far Are We From Believable AI Agents? A Framework for Evaluating the Believability of Human Behavior Simulation. <https://doi.org/10.48550/arXiv.2312.17115>
- [104] Hui Ye, Chufeng Xiao, Jiaye Leng, Pengfei Xu, and Hongbo Fu. 2025. Mo-GraphGPT: Creating Interactive Scenes Using Modular LLM and Graphical Control. <https://doi.org/10.48550/arXiv.2502.04983>
- [105] Fabio Massimo Zanzotto. 2019. Human-in-the-loop artificial intelligence. *Journal of Artificial Intelligence Research* 64 (2019), 243–252.
- [106] Runcong Zhao, Wenjia Zhang, Jiazheng Li, Lixing Zhu, Yanran Li, Yulan He, and Lin Gui. 2024. NarrativePlay: An Automated System for Crafting Visual Worlds in Novels for Role-Playing. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*. AAAI Press, Palo Alto, CA, Article 2866, 3 pages. <https://doi.org/10.1609/aaai.v38i21.30589>

A Participants' Background Information

Table 4: Pilot study participants' background information. Gender is abbreviated as M = Male and F = Female. Weekly Frequency is abbreviated as Wk. Freq. Experience is abbreviated as Exp.

ID	Gender	Age	Major	Everyday Storytelling		Interactive Vignettes	
				Viewing Wk. Freq.	Authoring Wk. Freq.	Playing Wk. Freq.	Authoring Exp.
P1	M	21	Brain Science	>5	3–5	1–2	None
P2	M	26	Electrical Engineering	>5	1–2	>5	Built storyworlds in Minecraft
P3	F	23	Culture Technology	>5	3–5	1–2	None
P4	M	35	Material	>5	3–5	3–5	None
P5	F	25	Industrial Design	>5	>5	>5	Tried to make a 3D character in Unity
P6	F	27	Chemistry	3–5	>5	1–2	None
P7	F	22	Biomedical Sciences	>5	>5	3–5	None
P8	M	21	Computer Science	>5	3–5	0	None
P9	M	32	Computer Science	1–2	>5	3–5	None
P10	F	21	Education	3–5	1–2	3–5	None
P11	M	19	Robotics	3–5	3–5	>5	None
P12	M	24	Transportation Engineering	>5	>5	1–2	None
P13	F	23	Electrical Engineering	>5	3–5	3–5	None

Table 5: User study participants' background information. Gender is abbreviated as M = Male and F = Female. Icons in the last column indicate experience categories: 🗨️ = Text Generation, 🖼️ = Image Creation, 🎥 = Video Creation, and </> = Coding. Weekly Frequency is abbreviated as Wk. Freq. Experience is abbreviated as Exp.

ID	Gender	Age	Major	Everyday Storytelling			Interactive Vignettes		AI-Assisted Authoring Exp.
				Viewing Wk. Freq.	Authoring Wk. Freq.	Wk.	Playing Wk. Freq.	Authoring Exp.	
P1	M	22	Aerospace Engineering	>5	>5		1–2	None	🗨️
P2	M	26	Cognitive Sciences	3–5	3–5		1–2	None	🗨️, 🖼️
P3	M	21	Computer Science	>5	3–5		>5	None	🗨️, </>
P4	F	21	Business and Technology Management	>5	3–5		3–5	None	🗨️
P5	M	21	Industrial Design	3–5	1–2		>5	None	🗨️, 🎥
P6	M	26	Culture Technology	3–5	>5		1–2	None	🗨️
P7	M	22	Bio and Brain Engineering	>5	3–5		3–5	None	🗨️, 🖼️
P8	M	22	Physics	>5	>5		0	None	🗨️
P9	F	25	Business and Technology Management	1–2	>5		3–5	None	🗨️, 🖼️
P10	M	21	Computer Science	3–5	3–5		3–5	None	🗨️, </>
P11	F	19	Undeclared Major	>5	>5		>5	Built an RPG using Unity	🗨️
P12	M	23	Computer Science	>5	1–2		1–2	None	🗨️, </>
P13	M	18	Undeclared Major	>5	3–5		3–5	None	None
P14	M	23	Physics	>5	>5		1–2	None	🖼️
P15	M	22	Computer Science	3–5	1–2		0	None	🖼️, </>
P16	M	23	Computer Science	>5	3–5		1–2	None	</>