Dramamancer: Interactive Narratives with LLM-powered Storylets

Tiffany Wang
Midjourney
San Francisco, California, USA
twang@midjourney.com

Yuqian Sun
Storytelling Lab
Midjourney
London, United Kingdom
ysun@midjourney.com

Yi Wang
Storytelling Lab
Midjourney
San Francisco, California, USA
ywang@midjourney.com

Melissa Roemmele
Midjourney
San Francisco, California, USA
mroemmele@midjourney.com

John Joon Young Chung
Midjourney
San Francisco, California, USA
jchung@midjourney.com

Max Kreminski
Storytelling Lab
Midjourney
San Francisco, California, USA
mkreminski@midjourney.com

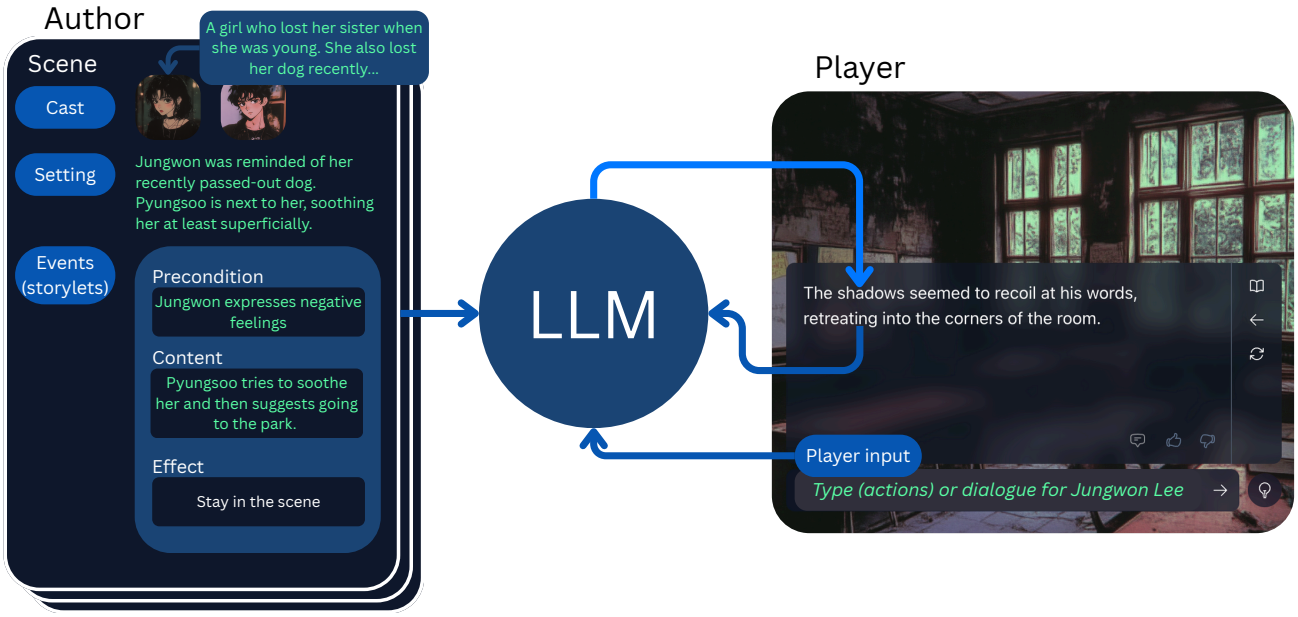

Figure 1: Dramamancer. The author user can structure stories by describing the setting and events (or storylets, which are composed of precondition, content, and effect) in natural language (green texts in `Author`). The player can experience the authored stories while making their character's actions in natural language (green texts in `Player`). LLMs power the generation of the story texts, conditioned by the author-provided story configuration, the player input, and the previously generated story.

Abstract

We present Dramamancer, an interactive narrative (IN) system that bridges authorial intent and audience input with natural language-based interactions powered by large language models (LLMs). Authors describe story settings in natural language and configure events via storylets, with natural language preconditions and contents. Players then experience these authored stories by inputting character actions in natural language, while the system adaptively generates story texts. LLMs serve as a narrative engine, dynamically incorporating both the authors' intent and the player's actions to flexibly generate stories at their intersections. By integrating LLMs with storylet-based IN authoring, Dramamancer aims to lower the authorial burden of creating INs while facilitating flexible player responses beyond predefined choices. We hope this demonstration inspires the UIST community to envision new types of interactive media that leverage intelligent technologies to bridge authors and audiences.

UIST Adjunct '25, Busan, Republic of Korea

ACM ISBN 979-8-4007-2036-9/25/09
https://doi.org/10.1145/3746058.3758995

CCS Concepts

• **Human-centered computing** → **Natural language interfaces**;
• **Computing methodologies** → **Natural language generation**.

Keywords

Interactive Narratives, AI-bridging, large language models

ACM Reference Format:
Tiffany Wang, Yuqian Sun, Yi Wang, Melissa Roemmele, John Joon Young Chung, and Max Kreminski. 2025. Dramamancer: Interactive Narratives with LLM-powered Storylets. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST Adjunct '25), September 28–October 01, 2025, Busan, Republic of Korea.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3746058.3758995

1 Introduction

Interactive Narratives (IN) engage audiences as active participants in authored stories by responding meaningfully to their actions within the story world. A core challenge in designing INs involves balancing narrative coherence, authorial intent, and rich player responsiveness to create immersive audience experiences [8].

Conventional IN authoring tools employ various constructs for authoring content in response to anticipated player actions. One popular construct is the *storylet* [5]—discrete, self-contained narrative units comprising of *preconditions* (triggers that determine when the storylet activates), *content* (narratives presented to the player once activated), and potentially *effects* (updates to the game state once activated). The modularity of storylets allows authors to gradually expand player options and reorganize narratives more easily compared to linear or branching counterparts. However, they still impose a significant authorial burden [2, 7] in that authors must craft sophisticated logical preconditions matching complex game state combinations [3], and player actions remain restricted to predefined options.

Large language models (LLMs) significantly improve player responsiveness by enabling free-form natural language input and just-in-time narrative generation with minimal authorial burden [1]. However, LLMs distance authors from the final narrative content presented to audiences, making it difficult to maintain story coherence and enforce authorial intent.

This work introduces Dramamancer, a tool that balances player responsiveness and authorial control while maintaining low authorial burden. We achieve this by combining the modular, flexible authoring paradigm of storylets with the open-ended capabilities of LLMs. Specifically, Dramamancer allows authors to express their intent behind characters, settings, and storylets in natural language (green text in `Author` section of Figure 1) while players can perform actions in natural language when experiencing authored INs (green text in `Player` section of Figure 1). The LLM serves as a bridge [4, 6]: it receives both author intent and player inputs in natural language, then generates story content at their intersections, achieving player responsiveness with low authorial burdens. To date, our internal deployment of Dramamancer has gathered a library of 15 interactive stories by 8 distinct authors, including a handful of stories created as templates for easy modification by later authors. We hope that this demonstration will spark discussions around new types of interactive and intelligent media that can balance authorial intent and audience responsiveness.

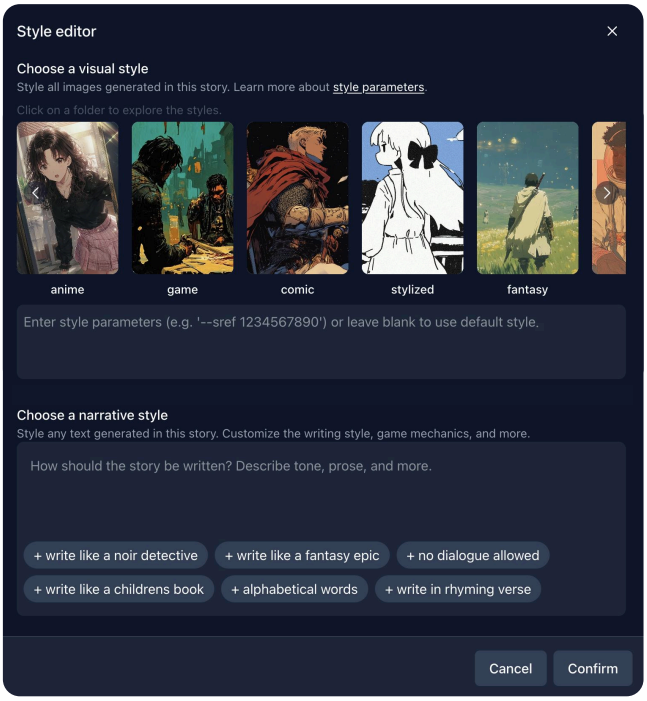

Figure 2: The author can specify their story's visual and textual styles with the style editor.

2 Dramamancer: Interaction

Dramamancer provides two main interfaces: an authoring interface for creating stories and a player interface for experiencing them.

2.1 Author Interactions

Authors can define four aspects of their story: style, characters, scenes and events.

2.1.1 Style. The authors first define the visual and textual styles of their story (Figure 2). Visual styles are used to generate character profiles, scene backgrounds, and event images, while textual styles are used to generate story texts. Dramamancer provides a predefined selection of categorized visual styles (e.g., `anime` unlocks `cinematic anime` or `classic anime`) with example images. Authors can also use custom parameters or leave styles blank. Textual styles are defined more flexibly with natural language, with presets available as shown in Figure 2.

2.1.2 Characters. For each character, authors can define a name, description (personality, background, relationships), and profile image (Figure 3). Names and descriptions are written in natural language. The first character serves as the player character.

Image Generation. Throughout Dramamancer, authors can generate images by writing natural language prompts to produce four

Figure 3: For each character, the author can specify their name, description, and profile image.

candidate options, or drag in images from Midjourney[1]. This applies to character profiles, scene backgrounds, and event images.

2.1.3 Scenes. Stories comprise multiple `scenes` (chapters) and each `scene` contains multiple `events` (storylets). For each scene, authors establish the `cast` (the characters involved), `setting` (location and context), `opening line` (the initial text that begins the scene), and background image.

2.1.4 Events. Events follow the storylet framework with *preconditions* ("If this condition is met") and *content* ("This happens") specified in natural language. Authors select *effects* ("Then go to") that transition to another scene or end the game, and can generate pop-up images for when events trigger (The bottom half of Figure 4 and Figure 5).

Dramamancer supports two event types:

- **Action-based events**: Triggered by natural language preconditions (Figure 4)
- **Time-based events**: Triggered after a specified number of turns, ensuring story progress even when players don't trigger story-advancing events (Figure 7)

Authors can receive LLM-generated suggestions for natural language preconditions by clicking on the bulb button at the top right of `Event` panel (Figure 4).

2.2 Player Interactions

Players experience stories beginning at the author-specified starting scene. The system generates narration and non-player character dialogues according to the author's specifications (the top of Figure 8). When prompted, players input actions or dialogue for their character using natural language (the bottom of Figure 8). Players can receive suggestions from the narrator on possible actions to take by clicking the bulb button at the bottom right of the interface. As players trigger events, they may see pop-up images pre-defined by the author (Figure 6) and transition between scenes or reach story endings.

[1]https://midjourney.com/

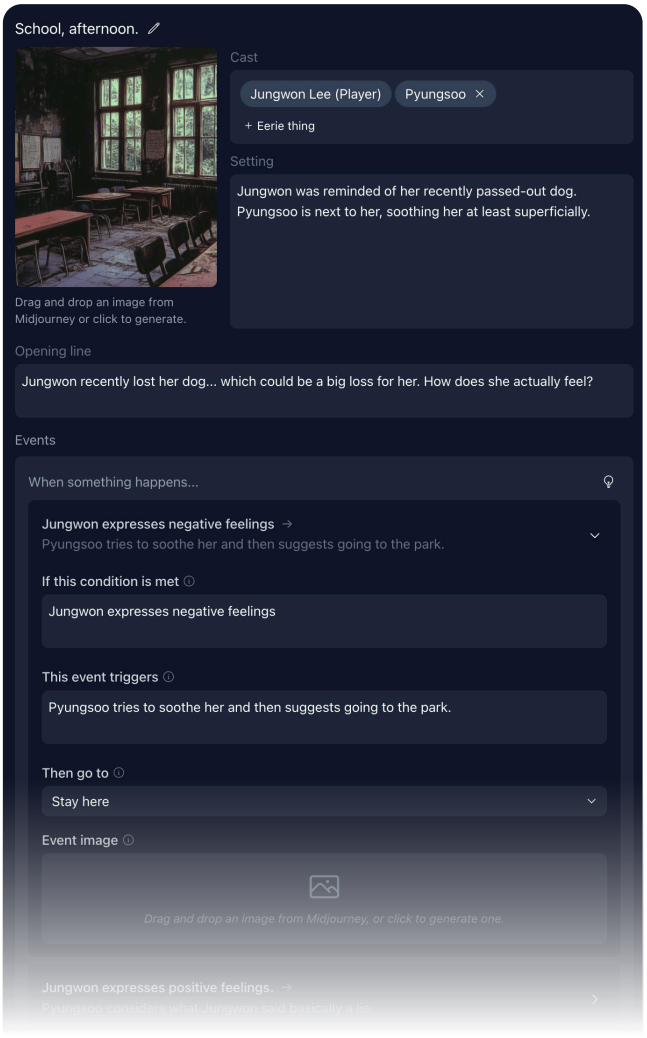

Figure 4: While a story can have multiple scenes, for each scene, the author can define the `cast`, the `setting`, the `opening line`, and `events`. Each `event` is formulated as a storylet.

3 Conclusion

We present Dramamancer, an IN system that bridges story authors and players through LLM-powered storylets. By allowing natural language to be used both in the writing of story events and in the expression of player actions, the system generates narrative content that balances authorial intent and player agency. With these LLM-powered storylets, Dramamancer addresses the long-standing challenge in IN systems of reducing authorial burden while maintaining player responsiveness. We hope that this demonstration sheds light on how we can design intelligent and interactive media that can connect with authors and audiences in novel ways.

Acknowledgments

We want to thank Midjourney for supporting this work.

Figure 5: The author can generate and define the image for event, which will pop up when the player triggers this event (Figure 6).

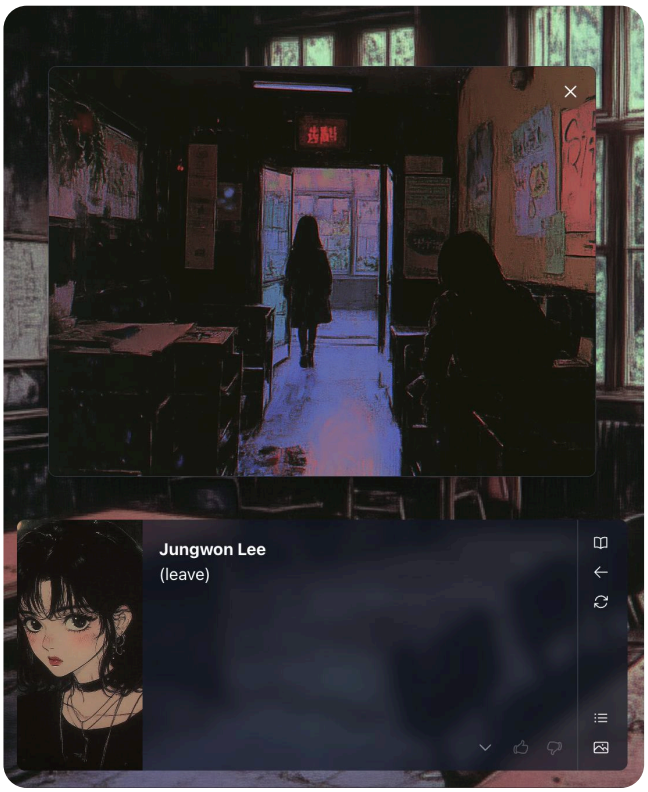

Figure 6: Event image pops up after the event for "Jungwon decides to leave" triggers.

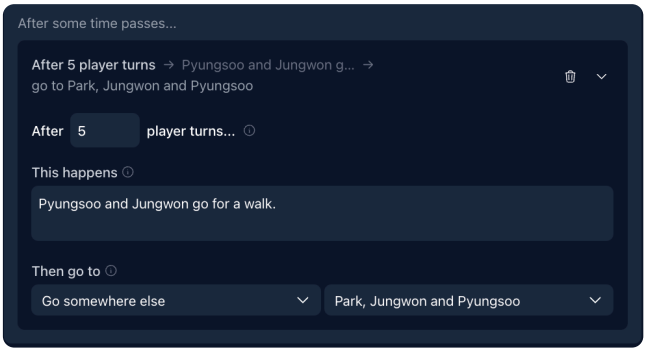

Figure 7: The `time-based event` triggers when a certain number of turns passes. It proceeds the story even when the player does not hit meaningful storylets to advance the story.

Figure 8: Player interface. For lines for non-playable characters or narration, the user can proceed by clicking the dialogue box (top). When Dramamancer prompts the player to make an action for the playable character, they can write down the action with natural language (bottom).

References

[1] Minh Hua and Rita Raley. 2020. Playing With Unicorns: AI Dungeon and Citizen NLP. *DHQ: Digital Humanities Quarterly* 14, 4 (2020).

[2] Joey Donald Jones. 2022. *Authorial Burden.* Springer International Publishing, Cham, 47–63. doi:10.1007/978-3-031-05214-9_4

[3] Joey Donald Jones and David Millard. 2024. Experiencing The Authorial Burden. In *Proceedings of the 35th ACM Conference on Hypertext and Social Media* (Poznan, Poland) *(HT '24)*. Association for Computing Machinery, New York, NY, USA, 78–87. doi:10.1145/3648188.3675134

[4] Taewook Kim, Hyomin Han, Eytan Adar, Matthew Kay, and John Joon Young Chung. 2024. Authors' Values and Attitudes Towards AI-bridged Scalable Personalization of Creative Language Arts. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 31, 16 pages. doi:10.1145/3613904.3642529

[5] Max Kreminski and Noah Wardrip-Fruin. 2018. Sketching a Map of the Storylets Design Space. In *Interactive Storytelling: 11th International Conference on Interactive Digital Storytelling, ICIDS 2018, Dublin, Ireland, December 5–8, 2018, Proceedings* (Dublin, Ireland). Springer-Verlag, Berlin, Heidelberg, 160–164. doi:10.1007/978-3-030-04028-4_14

[6] 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 (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 333, 18 pages. doi:10.1145/3706598.3713363

[7] Stacey Mason. 2021. *Responsiveness in narrative systems.* University of California, Santa Cruz.

[8] Mark Owen Riedl and Vadim Bulitko. 2013. Interactive narrative: An intelligent systems approach. *Ai Magazine* 34, 1 (2013), 67–67.