An Authoring Framework for LLM-Based Drama Managers

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

Yuqian Sun
ysun@midjourney.com
Midjourney
San Francisco, California, USA

Yi Wang
ywang485@gmail.com
Independent Researcher
Richmond, California, USA

Melissa Roemmele
melissa@roemmele.io
Independent Researcher
San Francisco, California, USA

Shm Garanganao Almeda
shm.almeda@berkeley.edu
UC Berkeley
Berkeley, California, USA

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

Max Kreminski
mkremins@cornell.edu
Cornell Tech
New York, New York, USA

Figure 1: Dramamancer scenarios are presented as visual novels that accept open-ended text input from the player, who speaks and acts as a specific player character. An LLM-based drama manager "narrates" everything else that happens in the storyworld, displaying non-player characters and generating dialogue for these characters at appropriate moments. When specific natural language conditions specified in advance by the scenario's author are met (e.g., if the player does something that should progress the scenario to the next act), a corresponding human-authored line of instructions is inserted into the LLM's context window. Such instructions alter the drama manager's behavior and may trigger actions such as character entrances and exits, scene transitions, or the playback of prescripted dialogue to advance the plot.

This work is licensed under a Creative Commons Attribution 4.0 International License.
I3D Companion '26, San Francisco, CA, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2667-5/2026/05
https://doi.org/10.1145/3807895.3807937

Abstract

One approach to interactive storytelling involves the construction of a *drama manager*: an intelligent agent that can maintain a storyline's coherence, even in the presence of unpredictable player actions, by improvising appropriate responses to whatever the player

says or does. LLM-based drama managers are able to respond to a wide range of open-ended player inputs, but on their own may struggle with pacing, structure, and narrative progression. We implement a dynamic context engineering framework that allows for the control of an LLM-based drama manager via *storylets*, which alter the LLM's context in systematic ways when certain narrative conditions are met. Our framework, Dramamancer, makes it easier to keep LLM-based drama managers "on track" from a human author's perspective without sacrificing the capacity to handle a wide range of potential player inputs.

CCS Concepts

• **Software and its engineering** → **Interactive games**; • **Applied computing** → **Media arts**; • **Human-centered computing** → *Natural language interfaces*; • **Computing methodologies** → *Natural language generation.*

ACM Reference Format:
Tiffany Wang, Yuqian Sun, Yi Wang, Melissa Roemmele, Shm Garanganao Almeda, John Joon Young Chung, and Max Kreminski. 2026. An Authoring Framework for LLM-Based Drama Managers. In *Companion Proceedings of the Symposium on Interactive 3D Graphics and Games (I3D Companion '26), May 13–15, 2026, San Francisco, CA, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3807895.3807937

1 Dramamancer

The medium of *interactive drama*—epitomized by experiences like *Façade* [Mateas and Stern 2003], which allow the player to speak and act freely as a character in an unfolding dramatic situation—attempts to bridge open-ended roleplay with a sense of participation in a tightly scripted and compelling plot. The author of an interactive drama scenario, however, faces the challenge of constructing a scenario-suitable *drama manager* [Chen et al. 2009; Roberts and Isbell 2008]: an intelligent agent that can respond appropriately to a wide range of unexpected player inputs, keeping the story "on the rails" from the author's perspective without excessively restricting the player's perceived freedom of action. The drama manager must puppet the non-player characters involved in the dramatic scenario, generate appropriate dialogue for these characters to speak, initiate scene or act changes when appropriate, and generally take on a wide range of responsibilities related to maintaining the perceived coherence and momentum of the narrative scenario—all without the author present to supervise the agent's actions at runtime.

In many ways, large language models (LLMs) seem well suited to drama management; their out-of-the-box ability to respond to open-ended textual inputs makes them especially versatile compared to older computational approaches to drama management, such as story planning [Riedl and Young 2010]. Researchers have employed LLM-based agents as dungeon masters in tabletop roleplay scenarios [Callison-Burch et al. 2022], and platforms like AI Dungeon [Hua and Raley 2020; Mitchell 2021] make successful use of LLMs as something between a writing partner and a narrator of non-player actions in an open-ended roleplay scenario. However, it is also well known that LLMs generally struggle to maintain important experiential qualities of story—such as pacing, tension, narrative structure [Tian et al. 2024], and a sense of consistent forward progression [Jeong and Lee 2025]—on their own.

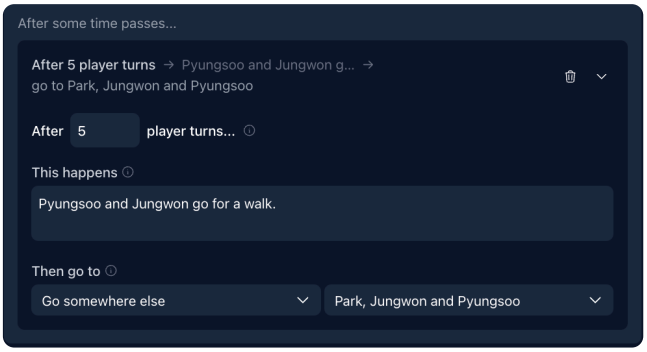

Figure 2: A *fallback storylet* triggers automatically after a certain number of turns have passed with no other storylets being triggered.

We therefore introduce Dramamancer [Wang et al. 2025b,c]: a context engineering framework [Mei et al. 2025] for the authoring of scenario-specific LLM-based drama managers, inspired by the interactive narrative concept of *storylets* [Kreminski and Wardrip-Fruin 2018]. A storylet in Dramamancer consists of a natural language *trigger condition*, a corresponding natural language *instruction*, and possibly zero or more *effects*. When the LLM judges that a storylet's trigger conditions have just been met (e.g., if the player does something that should progress the story to the next act), the associated instructions are inserted into the LLM's context window, and the associated effects (if any) are executed by the Dramamancer system. Instructions alter the drama manager's behavior, while effects may trigger characters to enter or leave the scene; change the scene, perhaps visibly swapping in a new background; advance the plot to the next act; or otherwise alter the dramatic situation. Some storylets can also be configured as *fallbacks*; if no other storylet is triggered for a fixed number of turns, a fallback storylet (Fig. 2) may automatically be triggered to keep the story moving forward, even in the absence of player action aligned with the scenario author's storytelling goals.

Authoring a drama manager in Dramamancer consists largely of determining what sorts of dramatic moments are important to systematically recognize and creating storylets that can respond appropriately when these sorts of moments occur—due either to player input, LLM-improvised responses to player input, or both. The *story schema* that defines a Dramamancer scenario consists of a set of storylets that can progress the plot in appropriate directions, plus a set of *characters* (each represented by a visual portrait and some natural language description); a set of *scenes* (each associated with a background and some optional scene introduction text); and a natural language description of the *style* in which the LLM should generally write when narrating and generating dialogue.

Dramamancer is open source[1] and has been used by dozens of authors to create a wide variety of different scenarios, many experimenting with new game mechanics that are only made possible by open-ended player input. The most similar extant framework, Orchid [Wu et al. 2025], has been developed in parallel with Dramamancer, and allows scenario authors to explicitly define a top-down

[1] https://github.com/dramamakers/dramamancer

state machine (with state transitions triggered by the recognition of open-ended narrative conditions) rather than a bottom-up collection of storylets to guide narrative progression; both frameworks represent approaches to bridging the open-endedness of LLM-based drama management with the tighter authorial control afforded by explicitly authored plot progression rules. As we continue to work with Dramamancer and refine its authoring affordances, we have also found ourselves developing new authoring tools meant to help authors anticipate and shape the range of playthroughs that players might experience in a particular scenario [Wang et al. 2025a]; we believe this represents a strong potential direction for future work.

Acknowledgments

Thanks to the Dramamakers team for helping us explore Dramamancer's authoring affordances more fully—and in particular to Hongwei Zhou, Nicholas Jennings, and Isabel Li.

References

Chris Callison-Burch, Gaurav Singh Tomar, Lara J Martin, Daphne Ippolito, Suma Bailis, and David Reitter. 2022. Dungeons and Dragons as a dialog challenge for artificial intelligence. *arXiv preprint arXiv:2210.07109* (2022).

Sherol Chen, Mark Nelson, and Michael Mateas. 2009. Evaluating the authorial leverage of drama management. In *Proceedings of the AAAI conference on artificial intelligence and interactive digital entertainment*, Vol. 5. 136–141.

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

Jin Jeong and Tak Yeon Lee. 2025. LIGS: developing an LLM-infused game system for emergent narrative. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*.

Max Kreminski and Noah Wardrip-Fruin. 2018. Sketching a map of the storylets design space. In *International Conference on Interactive Digital Storytelling*. Springer, 160–164.

Michael Mateas and Andrew Stern. 2003. Façade: An experiment in building a fully-realized interactive drama. In *Game Developers Conference*, Vol. 2. 4–8.

Lingrui Mei, Jiayu Yao, Yuyao Ge, Yiwei Wang, Baolong Bi, Yujun Cai, Jiazhi Liu, Mingyu Li, Zhong-Zhi Li, Duzhen Zhang, et al. 2025. A survey of context engineering for large language models. *arXiv preprint arXiv:2507.13334* (2025).

Alex Mitchell. 2021. Repetition and defamiliarization in AI Dungeon and Project December. In *Proceedings of the 2021 Electronic Literature Organization Conference*.

Mark O Riedl and Robert Michael Young. 2010. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research* 39 (2010), 217–268.

David L Roberts and Charles L Isbell. 2008. A survey and qualitative analysis of recent advances in drama management. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning* 4, 2 (2008), 61–75.

Yufei Tian, Tenghao Huang, Miri Liu, Derek Jiang, Alexander Spangher, Muhao Chen, Jonathan May, and Nanyun Peng. 2024. Are large language models capable of generating human-level narratives?. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 17659–17681.

Tiffany Wang, Yuqian Sun, Yi Wang, Melissa Roemmele, John Joon Young Chung, and Max Kreminski. 2025b. Design Techniques for LLM-Powered Interactive Storytelling: A Case Study of the Dramamancer System. In *Wordplay: When Language Meets Games @ EMNLP 2025*.

Tiffany Wang, Yuqian Sun, Yi Wang, Melissa Roemmele, John Joon Young Chung, and Max Kreminski. 2025c. Dramamancer: Interactive Narratives with LLM-Powered Storylets. In *Adjunct Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*.

Yi Wang, John Joon Young Chung, Melissa Roemmele, Yuqian Sun, Tiffany Wang, Shm Garanganao Almeda, Brett A Halperin, Yuwen Lu, and Max Kreminski. 2025a. Elsewise: Authoring AI-Based Interactive Narrative with Possibility Space Visualization. *arXiv preprint arXiv:2601.15295* (2025).

Zhen Wu, Serkan Kumyol, Shing Yin Wong, Xiaozhu Hu, Xin Tong, and Tristan Braud. 2025. Orchid: A Creative Approach for Authoring LLM-Driven Interactive Narratives. In *Proceedings of the 2025 Conference on Creativity and Cognition*. 774–791.