

Loose Ends: A Mixed-Initiative Creative Interface for Playful Storytelling

Max Kreminski¹, Melanie Dickinson², Noah Wardrip-Fruin¹, Michael Mateas¹

¹University of California, Santa Cruz, ²Independent
{mkremins, nwardrip, mmateas}@ucsc.edu, meldckn@gmail.com

Abstract

We present Loose Ends, a mixed-initiative co-creative storytelling play experience in which a human player and an AI system work together to compose a story. Loose Ends specifically aims to provide computational support for managing multiple parallel plot threads and bringing these threads to satisfying conclusions—something that has proven difficult in past attempts to facilitate playful mixed-initiative storytelling. We describe the overall human-AI interaction loop in Loose Ends, including the implementation of the rules-based AI system that enables this interaction loop; discuss four examples of desirable mixed-initiative interactions that are possible in Loose Ends, but not in similar systems; and present results from a preliminary expert evaluation of Loose Ends. Altogether, we find that Loose Ends shows promise for creating a sense of coauthorship in the player while also mitigating the directionlessness reported by players of earlier systems.

Introduction

Mixed-initiative creative interfaces (MICIs) (Deterding et al. 2017; Liapis et al. 2016) aim to support a human user’s creativity by providing them with an artificially intelligent creative partner. In the domain of storytelling-oriented creative writing, most existing MICIs function by providing suggestions as to how a story might be continued, thereby injecting unexpectedness into the writing process (Calderwood et al. 2020) and providing an immediate answer to the question of “What happens next?” when the user would otherwise become creatively stuck (Kreminski and Martens 2022).

These existing MICIs have shown promise in several ways. In particular, MICIs that function by providing short-term story continuations have proven effective at suggesting viable next steps for a story (Roemmele and Gordon 2015); taking the story in unexpected directions (Kreminski et al. 2020a; Calderwood et al. 2020; Singh et al. 2022); and creating a sense of shared authorship (Samuel 2016) between the user and system (Kreminski et al. 2020a; Calderwood et al. 2020; Singh et al. 2022).

However, these existing MICIs also exhibit several recurring problems. Most prominently, because the continuations these systems provide take only local context into account,

they have a tendency to pull the story in unwanted directions (Roemmele and Gordon 2015; Calderwood et al. 2020; Singh et al. 2022) or to otherwise create a sense of long-term directionlessness (Kreminski et al. 2020a) that inhibits the development of coherent high-level story structure.

To address these problems, we created Loose Ends, a MICI for storytelling that aims to support the development of coherent longer-term story structure. By explicitly reasoning about multiple parallel plot threads and providing a mixed-initiative interface for managing long-term storytelling goals framed in terms of these plot threads, Loose Ends aims to provide suggestions that keep the story on track with respect to the development of character arcs, conflicts, and high-level narrative themes.

Our main contributions are:

- A co-creative AI system that can reason about threaded plot structure in relation to high-level storytelling goals, proactively suggest new goals based on past plot events, and suggest character actions that advance these goals
- An approachable user interface for interacting with this AI system to create stories
- A preliminary evaluation of our approach by five experts in computationally engaged storytelling, indicating that Loose Ends shows promise at mitigating directionlessness while preserving a sense of coauthorship

In addition to these contributions, we also make the current version of Loose Ends available to be played in a web browser¹ and release its codebase as open source².

Background

Loose Ends draws inspiration from several past attempts to facilitate playful mixed-initiative storytelling, particularly *Writing Buddy* (Samuel, Mateas, and Wardrip-Fruin 2016) and *Why Are We Like This?* (Kreminski et al. 2020a,b). Both of these systems allow players to specify storytelling goals that guide the direction of the running story by influencing what story continuations the system will suggest. Both systems generate continuation suggestions in the form of structured plot events rather than prose, using a rules-based

¹<https://itsprobablyfine.github.io/LooseEnds>

²<https://github.com/ItsProbablyFine/LooseEnds>

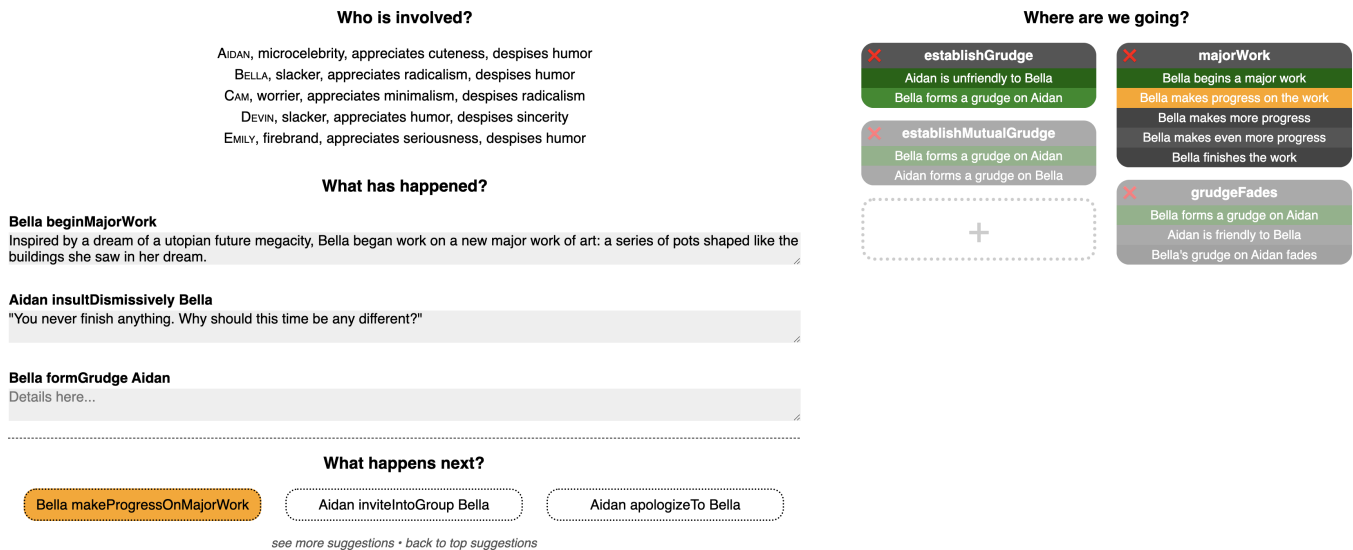


Figure 1: The Loose Ends user interface. The **Who is involved?** section displays basic information about a generated cast of five characters. The **What has happened?** section lists plot events that have taken place in the story so far, along with player-written text giving more details about these events. The **What happens next?** section shows AI-generated suggestions for what might happen next in the story. The **Where are we going?** section shows active storytelling goals, including transparent goals that have been suggested by the AI system rather than added by the player. One action suggestion (highlighted in orange in the bottom left) is being hovered over by the player; consequently, the impact this suggestion would have on the active storytelling goals if accepted (i.e. advancement of the `majorWork` goal) is also highlighted in orange on the right.

AI system rather than a language model to generate goal-relevant continuations. And both systems provide a story transcript that captures all past plot events in the form of a story outline, alongside player-written narration elaborating on the basic event descriptions generated by the system.

Loose Ends follows a similar architecture, although it differs from its predecessors in two key ways. First, its storytelling goals are more sophisticated than those in either predecessor system. Unlike in *Why Are We Like This?*, storytelling goals in Loose Ends specify *sequences* of events that must be added to the story for the goal to be satisfied (rather than individual events alone)—and unlike in *Writing Buddy*, storytelling goals in Loose Ends can be *parametrized* with specific characters and additional constraints. Second, the AI in Loose Ends is capable of suggesting new storytelling goals that are consistent with the story up until this point, rather than just steering action suggestions toward player-specified goals as in previous systems. Together, these changes result in a system that feels like an active writing partner while also guiding player-authored stories toward coherent longer-term structure.

Beyond plot event-based systems such as *Writing Buddy* and *Why Are We Like This?*, a number of attempts have also been made to facilitate mixed-initiative storytelling by providing continuation suggestions in the form of unstructured prose. Early examples of this approach can be found in the Say Anything (Swanson and Gordon 2012) and Creative Help (Roemmele and Gordon 2015) systems, which use case-based reasoning to find sentences similar to the user’s most recently typed sentence in a large database of preauthored stories, then suggest these sentences as continu-

ations. More recently, textual continuations provided by language models have been used to support storytelling in a relatively unmediated way (Manjavacas et al. 2017; Calderwood et al. 2020). Singh et al. (2022) finetune a large language model on a storytelling-relevant dataset and extend its completion suggestions to include images and sound as well as text, then evaluate this approach at scale. In each of these cases, purely text-based completions have been found to be pleasantly surprising and often relevant to the immediately previous parts of the story being told, but divergent from user-intended story structure in ways that require frequent revision by the user to maintain long-term direction.

One recent mixed-initiative storytelling support system that departs from the interaction paradigm of local continuation suggestion is TaleBrush (Chung et al. 2022), which instead aims to give users direct control of high-level story structure via the sketching of a visual fortune arc for the story’s main character. This approach has so far only been used to generate very short stories (on the order of five sentences long), and the coherence of the generated stories is limited, but this potential alternative means of specifying high-level storytelling goals still merits mention here. Another approach is that taken by MimiBrunnur (Stefnison and Thue 2018), which focuses on helping users create abstract story outlines that function as *generators* of stories rather than as the detailed backbone of a single story.

System Description

Loose Ends (Figure 1) is a mixed-initiative creative interface for playful storytelling. We specifically conceive of Loose

Ends as an AI-based *narrative instrument* (Kreminski and Mateas 2021): a system that can be played to produce narrative, in much the same way that a musical instrument can be played to produce music. The envisioned players of Loose Ends are the types of players who write retellings (Eladhari 2018) of their play experiences in games—especially emergent narrative games with sizable storytelling-oriented player communities, such as *The Sims* and *Dwarf Fortress*.

In the Loose Ends interaction loop, a human player repeatedly selects *action suggestions* furnished by the underlying AI system to continue the plot of a running story, using *storytelling goals* to steer the narrative toward player-desired long-term outcomes. Actions selected by the player are added to a running *story transcript*, and each action can be annotated with additional text by the player—for instance to narrate the action in greater detail.

Although Loose Ends as a system aims to be storyworld-agnostic, the version of Loose Ends presented here contains actions and storytelling goals that are specifically relevant to constructing stories about the development of character relationships and careers within a small community of artists. In the future, we envision that many different “playsets” for Loose Ends might be created, supporting the construction of stories set in many different kinds of storyworlds.

The AI system that powers Loose Ends consists of two major components. First is a **storytelling goals tracker** that updates a pool of active and possible storytelling goals as new plot events are added. Second is an **action suggestion generator** that generates and ranks potential suggestions for the next plot event in the story based on the currently active storytelling goals.

Storytelling Goals Tracker

Storytelling goals in Loose Ends are used to set and maintain the high-level direction of the story. Every goal is an instance of a *goal template*: a story sifting pattern written in the domain-specific logic programming language *Winnow* (Kreminski, Dickinson, and Mateas 2021). An example goal template is given in Appendix A.

A goal template describes a sequence of interrelated events that can be interpreted as satisfying a particular storytelling purpose or instantiating a particular kind of plot thread. For instance, the current version of Loose Ends includes templates for goals that introduce or develop character relationships (e.g., friendship or rivalry); internal conflicts (e.g., artistic or career struggles); and high-level narrative themes (e.g., moral themes related to the virtues of persistence in the face of adversity). There were 12 goal templates total in the version of Loose Ends evaluated here.

A *goal* is a partial match against a goal template, representing a sequence of past plot events that partially meet the goal template’s requirements. To *advance* a goal is to locate and accept an action suggestion that continues the sequence of events that match the underlying goal template. For instance, if a `majorWork` goal involving the character Aidan has been advanced past the first event (in which the goal’s main protagonist character begins work on a major art project) and a second event in which Aidan makes

progress on the project is added to the story, this goal will be advanced another step.

Goals can also be *cut off* if an event that violates one of the goal’s constraints is added to the story. For instance, if an `onARoll` goal involving the character Bella is active, but another character completes a major artwork before Bella manages to complete two major artworks in a row, this goal will be cut off, since a condition of the `onARoll` goal has now been violated.

Goals are *parametrized* by the characters that are involved in them, and multiple goals that are based on the same underlying goal template can be active concurrently as long as they pertain to a different configuration of characters. For instance, two `formGrudge` goals can be simultaneously active if either the character that holds the grudge, the target of the grudge, or both are different between the two goals. Additionally, if the player knows that they want a certain type of plot thread to be present in the story but does not know which characters they want to be involved, they can add a storytelling goal of the relevant type without any character parameters specified and allow the system to suggest possible ways of casting the available characters into this thread.

The Loose Ends user interface permits players to add goals manually (by selecting a goal template to instantiate as a goal, from a library of all available goal templates) and to remove goals that have already been established at any time. In addition, the AI system in Loose Ends constantly tracks and evaluates a pool of partial matches that the player has not established as goals. If one of these partial matches advances beyond a certain threshold (33% completion in the current version of Loose Ends), the system will automatically promote it to an active goal, rendered in a transparent style to indicate that this is a system-suggested goal rather than a player-added one. These goals can be removed by the player like any other (enabling the player to veto the system’s suggestions of additional storytelling goals), or the player can click on them to remove the transparency effect and notionally “lock them in” as player-intended goals.

Action Suggestion Generator

Action suggestions in Loose Ends are drawn from two pools of actions. The *basic actions pool* contains actions that are possible for any character at any time, regardless of social state, and remains fixed at all times. The *dynamic actions pool* is recalculated whenever a new event is added to the story, and contains actions that are only possible because of active storytelling goals that are in an appropriate state. For instance, when a complete `establishGrudge` goal between the characters Cam and Devin is active, the dynamic actions pool will contain actions that Cam can only take toward Devin because of their active grudge on Devin (such as sabotaging Devin’s most recent artwork). There were 32 action types total in the version of Loose Ends evaluated here: 20 basic actions and 12 dynamic actions.

Actions in general may be either *solo* (involving only a single character, the *actor* who takes the action) or *dyadic* (involving two characters, the actor who takes the action and the *target* toward whom the action is directed). Creating a minor artwork, for instance, is a solo action, while insulting

another character is a dyadic action. In addition, every action has an *event type* uniquely identifying the type of action that was performed and a list of zero or more *tags* that assign the action to high-level categories (such as *release* for actions in which the actor finishes and releases an artwork, *friendly* for actions in which the actor is friendly toward the target, and *harms* for actions that harm the target).

Action suggestions are recalculated every time the set of active storytelling goals changes. When calculating action suggestions, the action suggestion generator first iterates over all possible next actions (in both the basic and dynamic action pools) and determines, for each action, which storytelling goals would be impacted (either advanced or cut off) by the addition of this action to the story. Each action is then given a priority score, which is the sum of three factors:

- The number of active storytelling goals that this action would advance
- A constant factor (0.5) if this action is from the dynamic actions pool—i.e., if it is only possible because of an active storytelling goal
- A random factor (between 0 and 0.5) to randomly permute the priority of actions with the same base score

Actions are sorted by their score and displayed in order, with the three highest-scoring actions being pulled to the top of the action suggestions list. In this way, actions that relate most strongly to the active storytelling goals are prioritized for display, with randomness ensuring a degree of alternation between suggestions that advance parallel plot threads. When the user hovers over an action suggestion to consider it, the precalculated information about which storytelling goals this action would advance or cut off is used to display the ramifications of accepting this action in the storytelling goals pane on the right side of the user interface.

Interaction Examples

In conjunction, the Loose Ends AI and user interface permit several desirable interactions that are not possible in other mixed-initiative creative interfaces for storytelling. Four especially interesting examples of novel mixed-initiative interactions enabled by Loose Ends (all of which took place organically during evaluation) are presented below.

Discovering New Storytelling Goals Beyond simply suggesting action-level continuations to a running story in accordance with player-provided storytelling goals, Loose Ends can also infer new storytelling goals that are consistent with the story so far and proactively suggest these goals to the player. This often results in interactions where a player who would otherwise become uncertain of what to do next is inspired by, and begins pursuing, a system-discovered storytelling goal instead.

For instance, in Figure 2, the player has just completed two *establishGrudge* goals targeting the same character (Cam) have both been completed. At this point, Loose Ends automatically discovers and surfaces a successive character relationship development goal, in which Aidan and Bella (who both have grudges on Cam) bond over their shared dislike. The first two steps of this goal are already

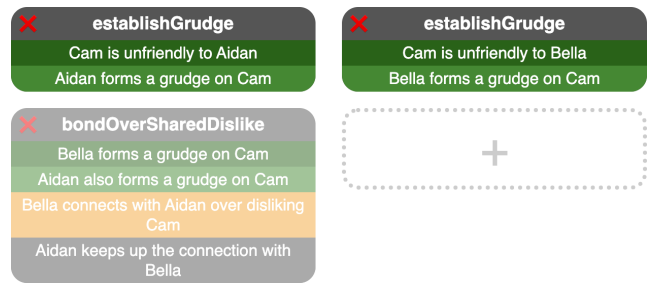


Figure 2: Based on events that were added to the story to complete two *establishGrudge* goals, Loose Ends has automatically discovered and surfaced a suggestion for another author goal (the *bondOverSharedDislike* goal) to spin off a new plot thread initiated by these events.




Figure 3: As the player considers an action that would advance one of their thematic goals but undermine another, the impact of the action on both thematic goals is highlighted, making the conflict apparent. The system knows that these goals are in conflict because *unless-event* clauses in the relevant goal templates mark these goals as cut off by characters achieving career success through alternative means.

complete, because the system has been tracking the possibility of surfacing this goal in the background, but it has only just now progressed far enough to be displayed.

Discovering Thematic Conflicts Loose Ends can make it apparent when a conflict has arisen between two active storytelling goals. For instance, in Figure 3, the player is simultaneously working toward two distinct thematic goals for the story and considering an action that will reward Emily with career success after she completes a major artwork. This would support the theme that persistent work on a single major project leads to success (*slowAndSteady*) but undermine the competing theme that the way to success is to create a rapid succession of more minor artworks (*quantityOverQuality*). When the impact of the considered action on all active author goals is visualized, the conflict between these goals is revealed to the player.

Resurfacing Dormant Plot Threads Because Loose Ends can maintain a larger set of active storytelling goals than the player can hold in their head all at once, action suggestions can serve to remind players of incomplete plot threads that they would otherwise forget to revisit. For instance, long-term storytelling goals like the *tryTryAgain* thematic goal (which requires a single character to repeatedly release artworks that are poorly received, before finally releasing one that is well-received) may temporarily fade into the background as the player focuses on another sub-

plot that weaves together a few distinct storytelling goals at once—but once this more pressing subplot is complete, actions advancing the earlier thematic goal will again rise to the top of the action suggestions pool, reminding the player to return to the previously initiated thread. This interaction pattern particularly helps to facilitate narrative *reincorporation* as discussed by Tomaszewski (2011).

Interleaving Parallel Plot Threads When multiple parallel plot threads are active and none of these threads has storytelling priority, the slight random permutation of equally ranked action suggestions means that Loose Ends by default tends to promote actions that alternately advance different threads. This can help players escape fixation (Gero 2011), in which they develop a narrow and premature focus on one plot thread or set of characters and forget about the possibility of developing others.

Evaluation Procedure

Since Loose Ends is still under active development, we conducted a preliminary and formative expert evaluation, intended to give us an initial sense of whether we have made qualitative progress toward our user experience goals. This evaluation was modeled on the evaluation of *Geminate* (Kreminski et al. 2020c), an earlier mixed-initiative co-creative system published at AIIDE.

We recruited five expert evaluators, all acquaintances of the first author, and all of whom are both experienced creative writers and researcher-practitioners in intelligent narrative technologies. Four of these evaluators (an assistant professor of computer science, two game industry narrative designers, and one independent creator of narrative games) hold a PhD in a relevant area, while the other (a PhD student in computational media) holds multiple relevant graduate degrees. All evaluators had past experience with mixed-initiative storytelling in general, and none had encountered Loose Ends before. Because our evaluators were familiar with the state of the art in mixed-initiative storytelling, they were readily able to compare Loose Ends to similar systems and judge what it does well or poorly in comparison.

Each evaluator participated in a single remote play session via Zoom. Each session was approximately one hour long and began with a brief (approximately 5-minute) introduction to the Loose Ends interface by one of the researchers. Subsequently, the evaluator constructed a single story using the Loose Ends interface while thinking aloud and sharing their screen. Once the story was complete, one of the researchers asked several unstructured interview questions to prompt reflection on play patterns they observed during the session. Both the think-out-loud and interview portions of the playtest sessions were recorded for later analysis. Finally, evaluators were administered a brief user experience questionnaire consisting of the following questions:

- Q1. What is your overall impression of the system?
- Q2. How easy was it to use the system?
- Q3. Were you able to use it without unnecessary effort?
- Q4. Did you feel a sense of control over the story?
- Q5. Was the system fun to use?

- Q6. Did you feel a sense of ownership of the story?
- Q7. Were you curious to see what would happen next in the story?
- Q8. Did you generally know what direction you wanted the story to go next?

Q1 was open-ended and qualitative, while Q2-Q8 were quantitative, with responses ranging from 1-5 (where 5 indicates the highest level of agreement with the premise of the question). Q1-Q5 were adapted directly from the *Geminate* expert evaluation questionnaire (Kreminski et al. 2020c), while Q6-Q8 were intended to elicit reflection on aspects of the co-creative storytelling experience that were frequently mentioned by playtesters of *Why Are We Like This?* (Kreminski et al. 2020a). A summary of evaluator responses to the quantitative questions is given in Table 1.

Question	E1	E2	E3	E4	E5	Avg
Q2. Usability	4	4	4	5	4	4.2
Q3. Effortlessness	4	5	5	5	5	4.8
Q4. Control	4	4	4	4	3	3.8
Q5. Fun	4	5	4	4	4	4.2
Q6. Ownership	3	4	3	3	3	3.2
Q7. Curiosity	4	5	4	3	4	4.0
Q8. Direction	4	5	4	4	3	4.0

Table 1: Summary of evaluators’ responses to quantitative survey questions. All responses were given on a numeric scale from 1-5, where 5 is highest agreement.

Evaluation Results

Directionlessness Is Mitigated. Our central design goal for Loose Ends was to mitigate the sense of high-level directionlessness reported by players during playtesting of *Why Are We Like This?* (Kreminski et al. 2020a) and assist in the development of stories that contain satisfying high-level structure. Both quantitative and qualitative evaluation responses suggest that Loose Ends successfully supports the development and maintenance of high-level narrative direction from the player’s perspective.

Quantitative survey responses related to sense of storytelling direction (Q8) ranged from 3-5, indicating that all evaluators had a sense of where they wanted the story next to go at a majority of points during the storytelling process. Additionally, all but one evaluator (E5) reported a score of 4 or higher in this category.

Qualitative think-out-loud remarks and interview responses are consistent with these quantitative results. In particular, two evaluators (E4 and E5) remarked unprompted on how they never experienced writer’s block or a sense of being stuck during the play process. Additionally, no evaluators explicitly reported a sense of aimlessness or insufficient medium-term direction at any point during their playthrough. This stands in stark contrast to the prevalence of self-reported aimlessness during playtesting of *Why Are We Like This?*, wherein four of five playtesters reported a sense of directionlessness at least once during play (Kreminski et al. 2020a).

Coauthorship Is Preserved. One open question for Loose Ends was whether the AI system could successfully preserve the sense of shared authorship that players experience in *Why Are We Like This?* while intervening more proactively in the storytelling process—including through the suggestion of new high-level storytelling goals. Both quantitative and qualitative evaluation responses suggest that Loose Ends succeeds in this regard.

Quantitative responses regarding sense of control over the story (Q4), sense of ownership of the story (Q6), and sense of curiosity regarding what would happen next in the story (Q7) are especially salient here. For control, all evaluators reported a score of at least 3 (indicating a moderate sense of control), and all but one (E5) reported a score of 4 (indicating a strong, but not complete, sense of control). For ownership, all evaluators reported a score of at least 3 (indicating a moderate sense of ownership), and one (E2) reported a score of 4 (indicating a strong, but not complete, sense of ownership). For curiosity, scores were distributed across the 3-5 range, indicating that all evaluators felt at least moderate curiosity, while all but one (E4) experienced either strong or very strong curiosity regarding the story's next direction. Taken together, these scores suggest that evaluators generally remained in control of the story while working with the system, but that they also created stories containing unexpected twists that they would be unlikely to invent if writing alone—to the extent that the AI system seemed to hold partial ownership of the stories that emerged.

Qualitative think-out-loud remarks and interview responses further support this interpretation. One evaluator (E5) felt that the play process reflected “a nice meeting in the middle” between player-led and system-led storytelling; another (E1) remarked that it “feels like the sweet spot for co-creativity”; and a third (E2) felt it to be a “good collaboration”: “kind of the dream” for mixed-initiative co-creativity.

Goal Alignment Is Unexpected and Fun. Four evaluators (E2-E5) remarked unprompted on how much they enjoyed it when the system correctly anticipated where they wanted the story to go next and offered options (especially storytelling goal suggestions) for continuing the story in a relevant direction. One evaluator (E2) was particularly pleasantly surprised by how often this took place during play. This suggests that the feeling of being seen or understood by the system can be a significant source of enjoyment during mixed-initiative storytelling, perhaps related to the aesthetic of responsiveness as described by Mason (2021).

Evaluators Found Loose Ends Easy to Use. Loose Ends was rated highly by evaluators on usability and (especially) lack of unnecessary effort involved in use, suggesting that it is considered highly usable in comparison to similar systems with which these evaluators were familiar. All evaluators reported a score of at least 4 for both Q2 (usability) and Q3 (effortlessness), and all but one evaluator (E1) reported a score of 5 for effortlessness, indicating unanimous agreement that Loose Ends is easy to use.

One caveat to this finding is that our evaluators, as experts in computationally engaged storytelling, were already familiar with several similar systems and used to putting up with

unpolished interfaces. Consequently, this finding might not generalize well to other player populations.

Some Players Want Prose-Level Suggestions. Evaluators used the freely editable text boxes in the story transcript in very different ways. Two evaluators (E1 and E4) mostly used them to write extended narration of high-level plot events, as we originally envisioned. One (E2) ignored the text boxes almost entirely. One (E3) used the text boxes to write short notes-to-self about why they chose certain actions from a storytelling perspective—a use-case we did not envision. And one (E5) initially used the text boxes to add terse narrative details for later expansion into full narration, but then stopped using them partway through play.

In qualitative think-out-loud remarks and interview responses, two evaluators (E1 and E2) both indicated that they wanted assistance in coming up with potential details for how certain high-level actions could have been narrated. E2 in particular (who made almost no use of the text boxes) stated that they would have found this additional narration-level support especially helpful.

Altogether, under the cognitive process model of writing (Flower and Hayes 1981; Gero et al. 2022), we find that Loose Ends currently provides assistance mostly at the planning stage, specifically in the creation of plot outlines. Expansion of support to later stages of the writing process represents a potential direction for future work.

Storyworld Inconsistencies Stand Out. The current version of Loose Ends makes use of a stateless, naïvely random action suggestion generator rather than a full-fledged social simulation to generate candidate action suggestions. Character relationship state is not tracked anywhere besides in storytelling goals related to friendship and rivalry, and most action types can be suggested between any pair of characters regardless of these characters' current relationship state. This leads to occasional generation of action suggestions that seem nonsensical from the perspective of a player who is tracking character relationship state mentally.

Three evaluators (E3-E5) commented at least once on this perceived occasional lack of consistency as a detriment to the overall storytelling experience in Loose Ends. This finding underscores the importance of storyworld consistency maintenance features for storytelling support—as suggested by several past studies, including Kreminski et al. (2019) and Calderwood et al. (2020). In the future, we intend to extend Loose Ends to use a more sophisticated suggestion generation mechanism that tracks substantially more character relationship state, hopefully alleviating this problem.

Common Feature Requests. Three evaluators (E1, E3, and E4) mentioned wanting to filter action suggestions to only display actions with particular characteristics, such as those of a particular event type or those involving particular characters. Three evaluators (E1, E3, and E4) mentioned a desire to express a temporary focus on a specific storytelling goal, so that the system would prioritize action suggestions that would advance this goal. Four evaluators (E1-E3 and E5) expressed a desire to minimize complete storytelling goals without removing them, in order to free up

more screenspace for incomplete goals. And finally, four evaluators (E1 and E3-E5) stated that they wanted more detailed information about a particular character’s traits or relationships to be immediately available while considering a suggested action involving that character. Going forward, we plan to add all of these features to Loose Ends in some form.

Evaluation Limitations. Our evaluation of Loose Ends is limited in several ways, particularly in terms of how evaluators were selected. Because evaluators were acquainted with the authors, some bias toward positive assessment of Loose Ends is likely; because evaluators were experts in mixed-initiative storytelling rather than novices, they likely found the system easier to use than novices might; and because the number of evaluators we employed is small, the quantitative results of evaluation might not generalize well. Comparison of our results to those from early-stage playtesting of *Why Are We Like This?* is still possible to some extent, since similar evaluators (interactive storytelling researchers who knew the system’s creators) were employed in *WAWLT* playtesting. However, a larger user study with players who are not researcher-level experts in mixed-initiative storytelling should be conducted in the future to determine more conclusively whether Loose Ends effectively supports storytelling among a more general audience.

Conclusion

Preliminary evaluation of Loose Ends, a novel mixed-initiative creative interface for storytelling, suggests that it may be able to preserve the desirable sense of coauthorship present in earlier systems while mitigating player-perceived narrative directionlessness. We hope that the formalization of jointly human- and machine-understandable storytelling goals presented here, and the idea of a mixed-initiative storytelling partner that can explicitly reason about and suggest high-level plot directions for a story (in addition to immediate continuations), will be taken up and further developed in the next generation of MICIs for storytelling support.

Appendix A: Example Storytelling Goal Template

Storytelling goal templates in Loose Ends are event sequence matchers written in the domain-specific logic programming language *Winnow*. Below is the code that implements the `bondOverSharedDislike` storytelling goal.

```
(pattern bondOverSharedDislike
  (event ?e1 where eventType: formGrudge,
    actor: ?c1, target: ?c3)
  (event ?e2 where eventType: formGrudge,
    actor: ?c2, target: ?c3)
  (event ?e3 where tag: friendly,
    actor: ?c1, target: ?c2)
  (event ?e4 where tag: friendly,
    actor: ?c2, target: ?c1)
  (unless-event where eventType: abandonGrudge,
    actor: ?c1, target: ?c3)
  (unless-event where eventType: abandonGrudge,
    actor: ?c2, target: ?c3))
```

Goals of this type match a sequence of four plot events (successively bound to the variables `?e1` through `?e4`), involving three characters (`?c1` through `?c3`), with arbitrarily many unrelated events interleaved between. Specifically, the goal will be satisfied if two characters `?c1` and `?c2` both form grudges on the same character `?c3`, and two reciprocally friendly interactions then happen between `?c1` and `?c2` (without either character first abandoning their grudge on `?c3`). For more information on the semantics of the *Winnow* language, see Kreminski, Dickinson, and Mateas (2021).

Appendix B: Example Output Story

The following text is an example story created by one of our evaluators. Each plot event is presented as a terse system-generated description in **bold**, followed by non-bold narration of the event written by the evaluator during play.

Aidan inviteIntoGroup Bella. Aiden meets Bella at an art opening, and invites her to a critique group.

Emily rejectSuperiority Bella. Emily hears Bella’s in Aidan’s cool critique group, and is envious, but says Bella’s not good enough to be in it.

Emily beginMajorWork. Emily starts work on a collage work that’s a thinly veiled-critique of Aidan and Bella’s group.

Bella formGrudge Emily. Bella hears what Emily’s new piece is about, and doesn’t like it.

Bella sendPostcard Aidan. Bella sends a postcard to Aidan about her new piece that’s been going through the critique group.

Aidan formFriendship Bella. Aidan and Bella become good friends through the group / shows / etc.

Emily receivePoorReview. Emily unveils her collage work, and it gets panned in the local art review, whose in the pocket of Big Aidan.

Emily askForHelp Bella. Emily asks Bella for some help winning over the local art critics.

Emily worryAboutMajorWork. Emily worries that Bella’s help won’t be enough to increase her reputation with the local art critics.

Emily makeProgressOnMajorWork. The collage grows to a series of collages about how messed up Aidan and Bella’s art group is.

Emily worryAboutMajorWork. Emily worries that despite the elaboration of the theme, it’s still too oblique to ring home for Aidan and Bella.

Bella formFriendship Aidan. Aidan and Bella become good friends through the group / shows / etc.

Emily makeProgressOnMajorWork. Emily continues making the collage series.

Emily makeProgressOnMajorWork. Emily continues making the collage series and secures a venue for the opening show for the series.

Emily finishMajorWork. She finishes the collages, mounts the show, and has the opening.

Emily receiveGoodReview. The show is a great hit, and the critics love it!

Bella shunFromGroup Emily. Bella reads the review, and finds out what the show was about, and says “you’ll never be part of our group”

Emily formGrudge Bella. Emily is like “you’re group is dumb, your art is facile, I don’t care.”

Bella beginMajorWork. Bella comes up with a collage OF HER OWN about Emily.

Cam buyLunchFor Devin. Meanwhile, the two artists to-tally outside the messed up politics of neo-collagists have lunch.

Devin apologizeTo Cam. Devin apologizes for not coming to Cam’s show...he was taking in this new collage form that just emerged...very experimental...very cool

Bella worryAboutMajorWork. Bella’s like “this collage stuff is actually really hard to make concrete statements about people with...damn is Emily actually a better artist than me??? Impossible!”

Bella makeProgressOnMajorWork. “Yeah impossible, this is ok...I guess....kinda”

Bella receiveNegativeFeedbackFrom Devin. Devin sees the work in progress and is like “are you biting on Emily’s style? This is derivative.”

Bella complainAboutMajorWork. Bella complains to Aidan about her collage piece, and that it’s more trouble than its worth.

Cam sendPostcard Devin. Cam sends a postcard with “wow, heard about the sick burn on Bella...haha”

Aidan apologizeTo Bella. “I’m sorry, I shouldn’t have introduced this conflict into your life...collage is a tempestuous medium, you should only approach it with pure intentions.”

Bella makeProgressOnMajorWork. Bella does.

Bella makeProgressOnMajorWork. Bella makes even more progress on the collage, but it’s becoming more about the travails and temptations of envy and jealousy and cliques.

Bella finishMajorWork. Shows the work, and invites Emily to see how she’s healed and moved on, and wants to bury the hatchet.

Bella receiveAward. Bella unexpectedly wins an award for her artwork, igniting a fresh, even more potent round of jealousy from Emily.

Acknowledgements

Max Kreminski conducted part of this research while in residence at Stochastic Labs.

References

Calderwood, A.; Qiu, V.; Gero, K. I.; and Chilton, L. B. 2020. How Novelists Use Generative Language Models: An Exploratory User Study. In *HAI-GEN + user2agent @ IUI*.
Chung, J. J. Y.; Kim, W.; Yoo, K. M.; Lee, H.; Adar, E.; and Chang, M. 2022. TaleBrush: Sketching Stories with Generative Pretrained Language Models. In *CHI Conference on Human Factors in Computing Systems*.
Deterding, S.; Hook, J.; Fiebrink, R.; Gillies, M.; Gow, J.; Akten, M.; Smith, G.; Liapis, A.; and Compton, K. 2017. Mixed-initiative creative interfaces. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, 628–635.
Eladhari, M. P. 2018. Re-tellings: the fourth layer of narrative as an instrument for critique. In *International Conference on Interactive Digital Storytelling*, 65–78. Springer.
Flower, L.; and Hayes, J. R. 1981. A cognitive process theory of writing. *College Composition and Communication*, 32(4): 365–387.

Gero, J. S. 2011. Fixation and commitment while designing and its measurement. *The Journal of Creative Behavior*, 45(2): 108–115.

Gero, K.; Calderwood, A.; Li, C.; and Chilton, L. 2022. A Design Space for Writing Support Tools Using a Cognitive Process Model of Writing. In *Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022)*, 11–24.

Kreminski, M.; Dickinson, M.; and Mateas, M. 2021. Winnow: A Domain-Specific Language for Incremental Story Sifting. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 17, 156–163.

Kreminski, M.; Dickinson, M.; Mateas, M.; and Wardrip-Fruin, N. 2020a. Why Are We Like This?: Exploring writing mechanics for an AI-augmented storytelling game. In *Proceedings of the 2020 Conference of the Electronic Literature Organization*.

Kreminski, M.; Dickinson, M.; Mateas, M.; and Wardrip-Fruin, N. 2020b. Why Are We Like This?: The AI architecture of a co-creative storytelling game. In *International Conference on the Foundations of Digital Games*.

Kreminski, M.; Dickinson, M.; Osborn, J.; Summerville, A.; Mateas, M.; and Wardrip-Fruin, N. 2020c. Germinate: A mixed-initiative casual creator for rhetorical games. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 16, 102–108.

Kreminski, M.; and Martens, C. 2022. Unmet Creativity Support Needs in Computationally Supported Creative Writing. In *Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022)*, 74–82.

Kreminski, M.; and Mateas, M. 2021. Toward Narrative Instruments. In *International Conference on Interactive Digital Storytelling*, 499–508. Springer.

Kreminski, M.; Samuel, B.; Melcer, E.; and Wardrip-Fruin, N. 2019. Evaluating AI-based games through retellings. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 15, 45–51.

Liapis, A.; Yannakakis, G. N.; Alexopoulos, C.; and Lopes, P. 2016. Can Computers Foster Human Users’ Creativity? Theory and Praxis of Mixed-Initiative Co-Creativity. *Digital Culture & Education (DCE)*, 8(2): 136–152.

Manjavacas, E.; Karsdorp, F.; Burtenshaw, B.; and Kestemont, M. 2017. Synthetic literature: Writing science fiction in a co-creative process. In *Proceedings of the Workshop on Computational Creativity in Natural Language Generation (CC-NLG 2017)*, 29–37.

Mason, S. 2021. *Responsiveness in Narrative Systems*. Ph.D. thesis, University of California, Santa Cruz.

Roemmele, M.; and Gordon, A. S. 2015. Creative Help: A story writing assistant. In *International Conference on Interactive Digital Storytelling*, 81–92. Springer.

Samuel, B. 2016. *Crafting Stories Through Play*. University of California, Santa Cruz.

Samuel, B.; Mateas, M.; and Wardrip-Fruin, N. 2016. The design of Writing Buddy: a mixed-initiative approach towards computational story collaboration. In *International Conference on Interactive Digital Storytelling*, 388–396. Springer.

Singh, N.; Bernal, G.; Savchenko, D.; and Glassman, E. L. 2022. Where to hide a stolen elephant: Leaps in creative writing with multimodal machine intelligence. *ACM Transactions on Computer-Human Interaction*.

Stefnisson, I.; and Thue, D. 2018. Mimisbrunnur: AI-assisted authoring for interactive storytelling. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 14, 236–242.

Swanson, R.; and Gordon, A. S. 2012. Say Anything: Using textual case-based reasoning to enable open-domain interactive storytelling. *ACM Transactions on Interactive Intelligent Systems (TiS)*, 2(3).

Tomaszewski, Z. 2011. On the use of reincorporation in interactive drama. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 7, 84–91.