Proceedings of the Eighteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment
(AIIDE 2022)

Open-World Narrative Generation to Answer Players' Questions

Cory Siler

Narrative Intelligence Lab, Department of Computer Science
University of Kentucky, Lexington, KY, USA 40506
cory.siler@uky.edu

Abstract

Planning-based narrative generation is effective at producing stories with a logically-sound flow of events, but it can be limiting due to the rigidity of its constraints and the high burden on the domain author to define story-world objects, initial states, and author and character goals. Giving the system the freedom to add objects and events to the story-world history arbitrarily can improve variety and reduce authorial burden, but risks leading to stories that seem jarringly contrived to the audience. I propose to use question-answering as the antidote to contrivance in a highly-generative interactive narrative system: By modeling the player's beliefs about the story world, inferring the implicit questions the player may be asking through their interactions, and answering those questions in a way consistent with the player's prior knowledge, a system could focus on creating cohesion in the ways that matter most to the player while accepting a degree of contrivance in the details that the player is likely to overlook.

Introduction

Planning-based methods for artificially intelligent narrative generation (Young 1999) excel at making stories well-structured; they can enforce the causal interconnectedness of events as well as the appearance of rational goal-directed character behavior (Riedl and Young 2010). However, many approaches struggle with generating an interesting variety of stories. One limitation of most planning-based narrative generators is that they require the domain author to fully define the objects in the story world, the characters' goals, and the world's initial state; while there been some work on *open world* generation where the system can modify the initial state (Riedl and Young 2005), giving an experience manager the flexibility human storytellers have to improvise new elements for the story is still an open problem. Another limitation is the rigidity of the character constraints in some narrative planning models; while it is valuable for characters to act as though they have beliefs, desires, and intentions, it may be unnecessary for a system to treat its agent model as a hard constraint on the same level as the laws of physics.

Suppose we relax some of the typical assumptions of planning-based story generation: Given a partial story structure with an incomplete set of states, world objects, events, and/or goals, we allow our generator to extend the structure arbitrarily as long as new additions are logically possible given the existing commitments. With few hard constraints on what is added, how do we guide generation to favor stories that are not overly contrived?

Human-authored fiction often accepts a certain degree of contrivance—someone happening to be in the right place at the right time to overhear a critical piece of information, a plan failing because of an oversight that a real person likely would not have made, a protagonist manifesting and then conquering a character flaw at the perfect moments to escalate and then resolve a dramatic situation, etc. Blatant "cheap tricks" (Ryan 2009) in the plot can be jarring, but audiences rarely demand that narratives (outside of certain genres such as rationalist fiction (McCain 2018)) be entirely free of contrivance; instead, a successful storyteller models the audience's mind to decide where justification and elaboration are needed and where they can be omitted.

In other words, the notion of whether a story "makes sense" does not exist in a vacuum; it requires a sensemaker, the audience, or in the case of interactive narrative in particular, the player. The player may imagine that there exists a ground-truth story world behind the scenes of what they observe, but the experience manager need not commit to one; it can instead choose details reactively to what the player does (Robertson and Young 2014). While prior works in player-centric planning and experience management often focus on what is *possible* based on a player's observations in a fully-enumerable world, my proposed work focuses on what is *plausible* in a player's mind when some of the details of an arbitrarily-extensible world are left to their imagination.

A player's actions within an interactive narrative can reveal implicit questions they have about the story (Ware 2017). I propose to use question-answering as the antidote to contrivance in a highly-generative interactive narrative system. When the player witnesses a non-player character take an action, or notices a certain item in a certain place, they may take the observation for granted, in which case the experience manager does not need to generate further information to explain the observation. Alternatively, the player may wonder what motivated the character to take the action, how the item got there, or what the eventual narrative role of either will be, in which case the experience manager must be prepared to create new material in response to the player's

investigations without "writing itself into a corner". I plan to develop methods to infer or anticipate the player's model of the story and their resulting questions, and to determine what content should be generated in response.

Past Work

My work so far has revolved around narrative planning, a form of plot generation that aims to balance character believability with authorial control over the story structure (Riedl and Young 2010). It does so by finding a plan (a sequence of actions representing story events, defined via logical preconditions and effects like in classical planning) that achieves an overall goal while also enforcing the appearance that each character acts upon their own individual goals.

The set of plans constituting solutions to a narrative planning problem can be seen as the intersection of two less-restrictive sets of plans: those that achieve the overall goal without limitations on character behavior, and those where all actions meet character constraints regardless of whether the overall goal is achieved. The former correspond to classical plans, while the latter resemble what might be produced by a truly multiagent system without central coordination. We analyzed several narrative planning problems in terms of the sizes of these less-restrictive sets compared to their intersection (Siler and Ware 2020); we put forth this form of analysis as a tool for choosing an appropriate architecture for particular narrative applications. We also hypothesized that a state-space narrative planner would be more efficient (in terms of number of search nodes explored) using a search strategy that prioritized overall goal achievement for problems where the set of overall-goal-achieving plans was smaller, and using a search strategy that prioritized character constraint satisfaction for problems where the set of character-constraint-satisfying plans was smaller; however, our results did not show a consistent pattern between relative set sizes and search strategy efficiency (Siler and Ware 2022).

We also introduced Sabre (Ware and Siler 2021), a system that finds solutions for narrative planning problems as defined by Shirvani, Farrell, and Ware (2018). Like its predecessor Glaive (Ware and Young 2014), Sabre is a forward-chaining state-space planner that finds an executed plan along with proofs that character actions conform to a model of intentionality. Unlike in Glaive, characters also have beliefs about the state of the world; unlike in nearly all prior narrative planners, characters can have beliefs about each other's beliefs, nested to an arbitrary depth. Any action taken by a character must be a nonredundant part of a hypothetical plan that starts in the state that the character believes the world is in, and ends with the character achieving one of their goals. A character c_1 can anticipate actions by another other character c_2 as part of c_1's own plan, but those actions must be consistent with c_1's beliefs about c_2's beliefs and goals.

Our benchmarking suggests that so far Sabre is fast enough to be useful for story graph generation (Ware et al. 2022) and similar offline applications, but not fast enough to embed directly in interactive narrative experiences where real-time responsiveness is critical. Furthermore, while experiments show that Sabre's nested-belief model can produce stories that agree with human readers' expectations more than a shallow-belief or no-belief model (Shirvani, Ware, and Farrell 2017), the psychology literature suggests that humans at any age do not consistently use their theory-of-mind ability to its full capacity (Keysar, Lin, and Barr 2003); as a result, past certain limits, Sabre may spend considerable time validating constraints that are ultimately lost on the player. My planned future work is inspired by these limitations: How can we let an experience manager reason about details that matter to the player while not wasting effort on details the player will not notice?

Planned Work

I propose interactive narrative experience management centered around a basic loop of *modeling* the story world and the player's beliefs about it, *inferring* the player's implicit questions, and *answering* questions through additions to the story world. This raises research questions about how the world and belief models should be structured, how inferences about questions can be made, and how answers can be chosen with the future trajectory of the interactive narrative in mind.

Modeling

A player-centric story knowledge representation needs at least two layers: one representing the immutable commitments that any future extensions to the story must respect—in particular, those entailed by observations given directly to the player—and one describing plausible worlds that the player might imagine existing outside of their observations based on what they know so far.

We say *plausible* rather than *possible* to distinguish this approach from traditional possible-worlds models. In these purely deductive models, an agent considers the set of all logically-possible worlds; as the agent observes new facts, worlds contradicting those facts are eliminated and the agent converges toward certainty about a single true world. Plausibility implies a form of nonmonotonic reasoning, such as Klassen, Levesque, and McIlraith's literary logic (2017). To paraphrase their example: By default, a player may assume dragons do not exist and rule out any worlds containing dragons. If the player encounters evidence of dragons, though, they may *add* dragon-containing worlds to their set of worlds under consideration. Furthermore, based on genre conventions, the player may only consider worlds that contain *fire-breathing* dragons, though encountering a non-fire-breathing dragon would be cause for another revision rather than a logical contradiction. While plausibility models have been studied in the context of story *understanding*, I am interested in adapting them for story generation.

Inferring

In an interactive narrative context, there has been extensive study of player modeling in terms of inferring what players hope to *do* in the story world, but less about what they

are trying to *learn*. These are not entirely separable concepts, though: I refer to Ram (1991)'s concept of *knowledge goals* and how they require an agent's existing knowledge and objectives to give them context. I propose to develop algorithms for determining players' knowledge goals based on their action histories and their other goals. Another possible useful tool is *narrative event salience* (Farrell, Ware, and Baker 2020; Flores and Thue 2017), which describes which events are likely to be memorable to a player and thus may highlight likely foci of player questions.

Answering

Answering questions about stories has prominently been studied (Graesser, Byrne, and Behrens 1992) and formulated in terms of computational plan structures (Cardona-Rivera et al. 2016; Farrell, Robertson, and Ware 2016) using the QUEST cognitive model (Graesser, Gordon, and Brainerd 1992). Answering an interactive narrative player's implicit questions, though, presents new challenges.

One challenge is that answers cannot always be stated directly to a player. For instance, suppose we have inferred that a player wants to know the goals that motivate a non-player character's action. Since most interactive narratives will not allow the player to peer directly into characters' minds, and it is not always desirable to require that the character's plans come fully to fruition (Ware and Young 2011), sometimes the system must find a way to show just enough of a character plan that the player can infer the rest.

Another challenge is that since we are considering an experience manager that can modify the story world's history to invent previously nonexistent answers, the consequences of those modifications may be limiting later on. For instance, after a long series of answers begetting new player questions, the experience manager may find itself in a position where the player is far off-track from the original scope of the story and it is difficult to achieve system goals. Balancing player exploration with more traditional desiderata of intelligent narrative such as preserving the domain author's envisioned story structure is an open area of investigation.

Related Work

My proposal shares with Castricato et al. (2021) the core loop of building a story iteratively by using question generation and answering to choose additions that support coherence. Their system, however, uses neural networks to analyze and output natural-language text that is most suitable for direct consumption by a human reader, while my approach will use symbolic reasoning to manipulate formal representations that are most suitable for interfacing with a game that interacts with a human player.

My proposal shares with Cardona-Rivera (2019) the metaphor of interactive narrative as a dialogue between player and experience manager, but with different assumptions about each. Cardona-Rivera's model assumes that the experience manager operates on a deterministic-state story world and has specific commitments about the player's future role in the story, and that the player prioritizes understanding and enacting the role meant for them. My model assumes a player who prioritizes understanding of the external story world, and that the underlying details of this world are nondeterministic and the experience manager is free to invent them based on the player's implicit queries.

My proposal shares with Robertson, Amos-Binks, and Young (2017) the concept of a nondeterministic-state story world, where what is possible is constrained by consistency with what the player has observed. Their approach handles a narrative planning problem where some variables may be left unassigned in the initial state, entailing multiple logically-possible worlds; as the player makes observations, the set of possible worlds shrinks, and the objective is to ensure the remaining worlds can eventually meet authorial goals despite possibly adversarial actions by the player. My model goes a step further into nondeterminism: By allowing new story-world elements to be introduced by the experience manager, it allows for a boundless set of *possible* worlds but instead tries to manage the set of *plausible* worlds that might be assumed based on what is known so far using a form of nonmonotonic reasoning. The set of plausible worlds may gain as well as lose elements when new information is revealed.

My proposal shares with Thue et al. (2017) the aim of increasing the range of stories that can be generated by giving the experience manager more freedom to introduce elements that are not preauthored. However, I diverge from their notion of initial and goal conditions; rather than requiring prespecified start and end points for the story, my proposed approach attempts to generate precisely as much story as needed needed to answer the player's implicit questions.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1911053. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

I thank the anonymous reviewers for their advice about both this paper and my work going forward.

References

Cardona-Rivera, R. E. 2019. *A model of interactive narrative affordances*. Ph.D. thesis, North Carolina State University.

Cardona-Rivera, R. E.; Price, T.; Winer, D.; and Young, R. M. 2016. Question answering in the context of stories generated by computers. *Advances in Cognitive Systems*, 4: 227–245.

Castricato, L.; Frazier, S.; Balloch, J.; and Riedl, M. 2021. Tell me a story like I'm five: story generation via question answering. In *Workshop on Narrative Understanding*.

Farrell, R.; Robertson, S.; and Ware, S. G. 2016. Asking hypothetical questions about stories using QUEST. In *International Conference on Interactive Digital Storytelling*, 136–146.

Farrell, R.; Ware, S. G.; and Baker, L. J. 2020. Manipulating narrative salience in interactive stories using Indexter's

Pairwise Event Salience Hypothesis. *IEEE Transactions on Games*, 12(1): 74–85.

Flores, L.; and Thue, D. 2017. Level of detail event generation. In *International Conference on Interactive Digital Storytelling*, 75–86. Springer.

Graesser, A. C.; Byrne, P. J.; and Behrens, M. L. 1992. Answering questions about information in databases. *Questions and Information Systems*, 229–252.

Graesser, A. C.; Gordon, S. E.; and Brainerd, L. E. 1992. QUEST: A model of question answering. *Computers & Mathematics with Applications*, 23(6): 733–745.

Keysar, B.; Lin, S.; and Barr, D. J. 2003. Limits on theory of mind use in adults. *Cognition*, 89(1): 25–41.

Klassen, T. Q.; Levesque, H. J.; and McIlraith, S. A. 2017. Towards representing what readers of fiction believe. In *International Symposium on Commonsense Reasoning*.

McCain, K. 2018. Sensibly organized: filling in gaps with *Harry Potter and the Methods of Rationality*. In *Prequels, Coquels and Sequels in Contemporary Anglophone Fiction*, 94–108. Routledge.

Ram, A. 1991. A theory of questions and question asking. *Journal of the Learning Sciences*, 1(3-4): 273–318.

Riedl, M. O.; and Young, R. M. 2005. Open-world planning for story generation. In *International Joint Conference on Artificial Intelligence*, 1719–1720.

Riedl, M. O.; and Young, R. M. 2010. Narrative planning: balancing plot and character. *Journal of Artificial Intelligence Research*, 39(1): 217–268.

Robertson, J.; Amos-Binks, A.; and Young, R. M. 2017. Directing intentional superposition manipulation. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 13, 232–239.

Robertson, J.; and Young, R. 2014. Finding Schrödinger's gun. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 10.

Ryan, M.-L. 2009. Cheap plot tricks, plot holes, and narrative design. *Narrative*, 17(1): 56–75.

Shirvani, A.; Farrell, R.; and Ware, S. G. 2018. Combining intentionality and belief: revisiting believable character plans. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 222–228.

Shirvani, A.; Ware, S. G.; and Farrell, R. 2017. A possible worlds model of belief for state-space narrative planning. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 101–107.

Siler, C.; and Ware, S. G. 2020. A good story is one in a million: solution density in narrative generation problems. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 123–129.

Siler, C.; and Ware, S. G. 2022. Solution density and search strategy in narrative generation. *IEEE Transactions on Games*. (forthcoming).

Thue, D.; Schiffel, S.; Gumundsson, T. .; Kristjánsson, G. F.; Eiríksson, K.; and Björnsson, M. V. 2017. Open world story generation for increased expressive range. In *International Conference on Interactive Digital Storytelling*, 313–316. Springer.

Ware, S. G. 2017. Mutual Implicit Question Answering for shared authorship: a pilot study on player expectations. In *Intelligent Narrative Technologies Workshop*, 259–265.

Ware, S. G.; Garcia, E. T.; Fisher, M.; Shirvani, A.; and Farrell, R. 2022. Multi-agent narrative experience management as story graph pruning. *IEEE Transactions on Games*. (forthcoming).

Ware, S. G.; and Siler, C. 2021. Sabre: a narrative planner supporting intention and deep theory of mind. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 99–106.

Ware, S. G.; and Young, R. M. 2011. CPOCL: a narrative planner supporting conflict. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 97–102.

Ware, S. G.; and Young, R. M. 2014. Glaive: a state-space narrative planner supporting intentionality and conflict. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 80–86.

Young, R. M. 1999. Notes on the use of plan structures in the creation of interactive plot. In *AAAI Fall Symposium on Narrative Intelligence*, 164–167.