Incorporating Authorial Intent into Generative Narrative Systems

Mark O. Riedl
School of Interactive Computing, Georgia Institute of Techology
85 Fifth Street NW, Atlanta, Georgia 30308, USA
riedl@cc.gatech.edu

Abstract

One of the major themes to emerge in interactive narrative research is *authorability* and *authorial intent*. With interactive narratives, the human author is not present at run-time. Thus authoring interactive narratives is often a process of anticipating user actions in different contexts and using computational mechanisms and data structures for responding to the participant. Generative approaches to interactive narrative, in which an automated narrative generation system assumes some of the authoring responsibility, further decouple the human designer from the participants experience. We describe a general mechanism, called *author goals*, which can be used by human authors to assert authorial intent over generative narrative systems.

Introduction

An interactive narrative is an approach to interactive entertainment in which a system attempts to tell a story to an interactive participant, such that the user is afforded the opportunity to make decisions that directly affect the direction and/or outcome of the story. One of the major themes to emerge in interactive narrative research is *authorability* and *authorial intent*. Authorial intent is the ability of an autonomous interactive system to reflect the intentions of the human designer – also called the human author. Because of participant agency, in an interactive narrative much of a participants actual run-time experience is influenced by the participants own actions. Unlike tabletop and live-action role-playing games, the human designer is not present at run-time and cannot make decisions about how the participants experience must be adapted to balance plot coherence and perceived participant self-agency. Thus authoring interactive narratives is often a process of anticipating user actions in different contexts and using computational mechanisms and data structures for responding to the user.

An approach to interactive narrative that may mitigate the authoring complexity is *generative drama management* (or *generative experience management* (Riedl et al. 2008) for non-dramatic contexts). The generative approach to interactive narrative suggests that if authoring branching stories is intractable for human authors then a computer system can generate story content in response to the actions and decisions of an interactive participant. The goal is to have the participants experience become part of an unfolding story. As the participant exerts his or her agency and deviates from the originally intended story, the experience manager invokes an automated story generation system to adapt, modify, or re-generate story content. Computers are very useful for performing recursive and repetitive tasks. A generative approach to narrative is favorable under circumstances in which there is too much variability for a human designer to foresee all eventualities.

Authoring interactive narrative content is thus a process of instilling a computational system with the ability to make the same decisions that the human designer would make in response to participant actions. That is, the human designers goal is to infuse his artistic vision and authorial intent into a computational system with the tools and data representations at hand. Generative experience management further decouples the human designer from the user. That is, not only will the human designer not be present to make decisions at run-time, but, with generative experience management, the human designer is also not responsible for authoring the narrative branches that will used to respond to user actions.

In this paper, we describe an attempt to provide mechanisms for injecting the designers intent into generative narrative systems. The mechanism we describe is called *author goals*, specialized data structures that are injected into planning-based generative systems that can be used by human authors to indicate preferences in narrative structure.

Related Work

We consider two types of generative narrative system. The first are narrative generation systems in which the purpose is to automatically produce a non-interactive narrative sequence. Many narrative generation systems (e.g., (Meehan 1976; Lebowitz 1985; Pérez y Pérez and Sharples 2001; Riedl and Young 2004)) are either based on planning or encapsulate planning-like processes. The second type of generative narrative system are those expressly developed to create real-time interactive experiences. See (Riedl et al. 2008) and (Roberts and Isbell 2008) for reviews of relevant work on interactive narrative systems. Note that two systems in particular, (Young et al. 2004) and (Riedl et al. 2008), explicitly rely on the recursive invocation of a narrative plan-

ner. By doing so, these systems build a tree where each child is a re-planned version of the parent narrative plan that handles some significant participant action.

Partial-order planning (c.f., (Weld 1994)) is a process of selecting and instantiating actions in a particular temporal order. Plans are comprised of actions. Actions have preconditions and effects. Preconditions dictate what must be true in the world for an action to be executed. Effects specify what will be different about the world once the action has been executed. Initially, the root of the search space is an empty plan and the goal state propositions are the only conditions that need to be satisfied. When an action (or the goal state) in a plan has a precondition that is not established by a preceding action (or the initial state) a new action is instantiated or an existing action is reused to satisfy that precondition or goal. An in-depth discussion of planning is beyond the scope of this paper.

Myers (1996; 2000) explored ways of controlling planners through advice through abstract specifications that are compiled into the planning domain. Advice includes task advice to identify goals and actions to be included, strategic advice to recommend how goals and actions are to be accomplished, and "evaluational" advice to inform the heuristic for overall plan evaluation. Thomas and Young (2006) extend Myers work by creating an environment for human authors to encode preferences through a domain metatheory.

Author Goals

Author goals serve two important purposes. First, author goals constrain the narrative search space such that it is impossible for a planner to produce a narrative that does not meet certain criteria imposed by the human author. Second, author goals can be used to force complexity in narrative generation.

Technically, author goals are a reformulation of *islands* for partial-order planners. Islands – a term coined to refer to a technique for controlling the form of solutions generated by planners (Hayes-Roth and Hayes-Roth 1979) – are intermediate states in a search space, through which all solutions to a planning problem must pass. In the early days of AI planning research, islands were used to inform the planner as to what valid solution plans should look like, conceptually speeding up the planning process. Potential solutions that do not satisfy each island state description at some point between the initial state and the end state are pruned.

Islands are tools for making state-space search practical for planning purposes. However, many modern planners use plan-space search. See (Weld 1994) for a discussion of the practical advantages of partial-order plan-space search over state-space search. As tools for making planning more pragmatic, islands are not typically necessary for partial-order plan-space search since they can search deeper ply. We use islands as a way for the human user to inject guidance into the narrative generation process and to force the planner to consider more complex action sequences.

Authorial Intent with Author Goals

A narrative generator necessarily operates without a human-in-the-loop. Author goals provide the ability to provide rough direction for what must occur within solutions. Further, a generative drama management system will automatically produce branches. Because those branches may or may not preserve elements from the root narrative the human author is encouraged to provide additional meta-data. The first meta-data is the outcome. The outcome is a description of the state of the story world after the story is complete. The second type of meta-data is author goals. In terms of authorial intent, an author goal indicates that there is a state of the world that must be achieved between the time the narrative starts and the outcome, and that any plan cannot be considered complete unless that world state is at least momentarily true.

Complexifying Narrative Plans

There are many narratives in which states reverse themselves one or more times. For example: a character that begins rich, becomes poor, and finally regains the state of being rich. These phenomena are challenging for planners without some form of guidance. For example, if a planner were given an initial state in which a character was rich and an outcome state in which the character is rich, the planner would simply indicate that there was no problem to solve. Author goals can be used to force the generator to consider substantially more complex plans in which some intermediate state, such as the character becoming poor, must be integrated into the resultant narrative structure.

Incorporating Author Goals into Planning

In our computational representation of narrative, author goals are implemented as a special type of plan step that have preconditions describing the intermediate world state but no effects. Author goals are provided at the time of planner initialization and describe world states that must be achieved at some intermediate time during plan execution. If more than one author goal is given, there can be pre-specified temporal links between them so that author goals must occur in the resulting, complete plan in a particular order. In this way, the existence of author goals constrains the space of plans that can be searched by the planner. That is, the planner cannot consider any plan in which the world state described by an author goal will not be achieved during plan execution.

To implement the ability for a planner to act on author goals, we make the following change to the way in which problems are described to the planner. Author goals are specified as sets of state propositions. State propositions, like a planning problem goal, define a set of states in which the given propositions are true. Author goals are specified in the form $(set_1 ... set_n)$ where each set_i is of the form $(authorgoal_1 ... authorgoal_N)$ and an $authorgoal_i$ is of the form $(proposition_1 ... proposition_n)$. Author goal sets are ordered, meaning that plans must achieve the author goals in the prescribed set order to be considered valid.

The above specifications for author-goals are translated into plan steps and inserted into the initial empty plan. For each author-goal a plan step data structure is created such that the state propositions make up the plan steps precondition list. The plan step has no effects. If there is ordering

```
(define (problem little-red)
 :inits ((character red) (human red)  (alive red)
         (character wolf) (monster wolf) (alive wolf)
         (character granny) (human granny) (alive granny)
         (character hunter) (human hunter) (alive hunter)
         (thing cake) (has red cake)
         (knows red granny) (knows granny red))
 :authorgoals ((((eaten red)))
               (((eaten granny)))))
 :outcome ((has granny cake)
           (:not (eaten red))
           (:not (eaten granny))))
```

Figure 1: The Little Red Riding Hood planning problem definition.

between author-goals, temporal links are added to the initial plan as well. Partial-order planning algorithms such as those based on (Weld 1994) do not need to be modified further. These planning algorithms treat the preconditions of the special author-goal plan steps as goals to be satisfied as normal. That is, the planner sees unsatisfied preconditions as flaws and attempts to instantiate an action (or select an existing action) that has an effect that unifies with the unsatisfied condition. The POP algorithm itself does not distinguish between an unsatisfied goal and an unsatisfied precondition on an existing action. Generative techniques not based on partial-order planner may require additional modifications to the generative algorithm itself to be able to take advantage of author goals.

Case Studies

In this section, we consider two case studies of generative narrative systems in which author goals were essential to their success. Both systems were built on a generative experience management framework described in (Riedl et al. 2008). The key consideration is that this framework uses partial-order planning technologies to generate narratives.

Little Red Riding Hood

The uses of author goals in a Little Red Riding Hood interactive narrative illustrate their necessity in complexifying narrative structure. Figure 1 shows the modified problem initialization in a PDDL-like language. An initial state defines characters, character traits, and relevant props and features of the world. The outcome is the goal: Granny has the cake and neither Little Red or Granny are in the state "eaten." The planning system is also initialized with an action library that describes ways in which the world can change. For example, characters can give things to other characters and some characters can eat other characters whole.

The author-goals define that a significant feature of generation is that Little Red and Granny should both, at some point, enter the state of being "eaten." Arguably, the Little Red Riding Hood domain could not be considered such without Little Red and Granny being eaten and in need of rescue. Note that the initialization parameters do not indicate how the author-goals or the outcome are achieved, only that they must be achieved. Author goals are necessary for

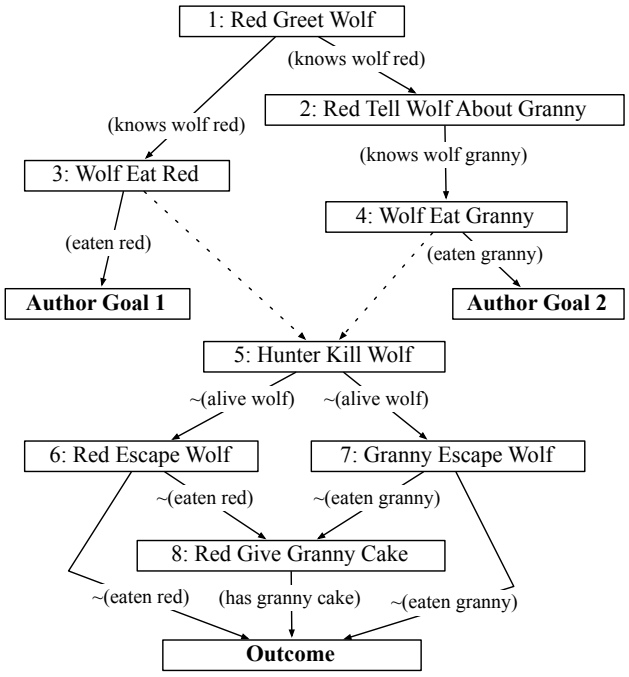

Figure 2: Example narrative plan set in the Little Red Riding Hood world.

complexification. Little Red and Granny begin the story in the state of being not eaten and end the story in the state of being not eaten. Without some indication that it is not desirable from a story-telling point of view that propositions about these characters states should change dramatically. Without author-goals, a planning algorithm could naively generate the following:

- Red gives Granny the cake
- The End.

The author goals given earlier prevent this by forcing the planner to consider substantially more complex plans in which Little Red and Granny become eaten and then are restored to not being eaten. See Figure 2 for one possible narrative plan that can be generated by respecting the author goals. Boxes are actions or author goals. Solid arrows are causal links – annotations on the plan that capture causal relationships. Dashed arrows are additional temporal constraints.

Socio-Cultural Awareness Training

The second case study is that of a socio-cultural awareness training prototype developed for the military. This system used an interactive narrative to expose trainees to socio-cultural situations in which dramatic situations unfold to challenge the trainee. In the scenario developed, two merchants in a foreign city are involved in a domestic dispute that escalates to violence, eventually involving the trainee acting in the role of peacekeeper. In the training scenario, author goals are used to express authorial intent. The purpose of the training scenario is to challenge the trainee by

creating dilemma situations where the appropriate course of action to take is not obvious without some deeper socio-cultural situational understanding.

Figure 3 shows the narrative plan. Some causal links are omitted for clarity. In this case study, there are three author goals: (a) the characters are established such that Hasan has acted suspiciously, Saleh has acted unfriendly, and Ali has acted unreliably, (b) a significant incident occurs such as an attack on the marketplace, and (c) the trainee is presented with two (on the surface level) equal possible courses of action, namely that Saleh is falsely accused of causing the incident and that Hasan is accurately accused by an unreliable character. The first author goal serves the purpose of introducing the characters. This set up stage is causally unnecessary to achieving the outcome, but is considered authorially important. The second author goal enforces the constraint that an incident occurs that instigates the final outcome dilemma since it could conceivably arise in other ways. Finally, the outcome state defines the dilemma conditions under which the trainee must act. The author goals were necessary because the authorial intentions were extremely hard to encode into the domain itself – the author's intentions were meta-constraints on the form of the physical action that actually occurs.

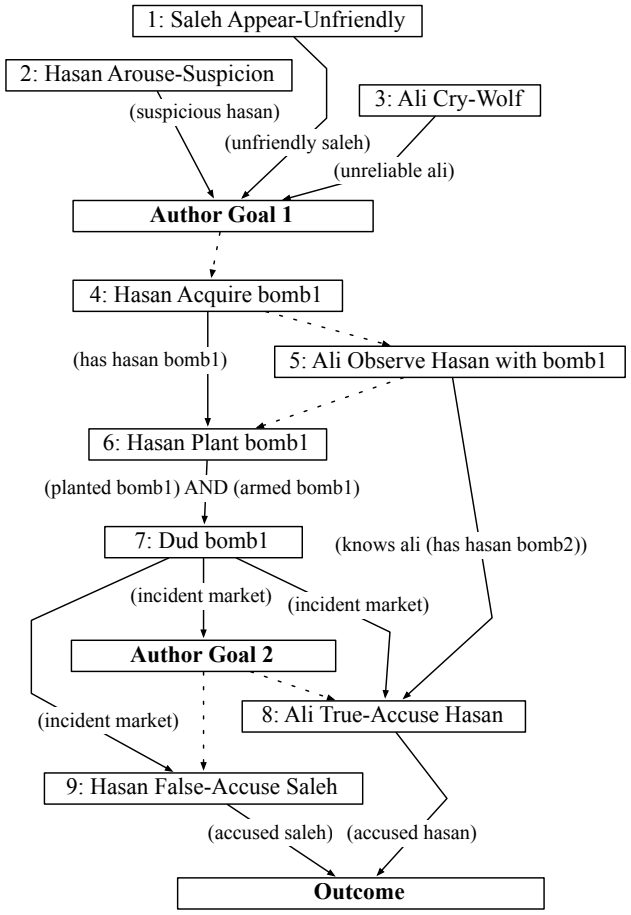

Figure 3: Example narrative for socio-cultural training.

Conclusions

The author goal mechanism described in this paper is an attempt to enable human authors to inject their preferences, intuitions, and requirements into the planning process. Author goals also have the pragmatic side effect that they can force more complexity in narrative plan solutions. In general, author goals constrain the planner to produce narrative plans with particular structures by pruning branches of the plan search space in which plans do not meet the author goals. We believe that enabling the author to inject control into the planning process will become more and more important as narrative systems such as narrative generators or interactive narratives acquire greater autonomy from the human author.

References

Hayes-Roth, B., and Hayes-Roth, F. 1979. A cognitive model of planning. *Cognitive Science* 3.

Lebowitz, M. 1985. Story-telling as planning and learning. *Poetics* 14.

Meehan, J. 1976. *The Metanovel: Writing Stories by Computer*. Ph.D. Dissertation, Yale University.

Myers, K. 1996. Advisable planning systems. In Tate, A., ed., *Advanced Planning Technology*. Menlo Park: AAAI Press.

Myers, K. 2000. Domain metatheories: Enabling user-centric planning. In *Proceedings of the AAAI Workshop on Representational Issues for Real-World Planning Systems*.

Pérez y Pérez, R., and Sharples, M. 2001. Mexica: A computer model of a cognitive account of creative writing. *Journal of Experimental and Theoretical Artificial Intelligence* 13.

Riedl, M., and Young, R. 2004. An Intent-Driven Planner for Multi-Agent Story Generation. In *Proc. of the 3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*.

Riedl, M.; Stern, A.; Dini, D. M.; and Alderman, J. M. 2008. Dynamic Experience Management in Virtual Worlds for Entertainment, Education, and Training. *International Transactions on System Science and Applications* 4(2).

Roberts, D., and Isbell, C. 2008. A survey and qualitative analysis of recent advances in drama management. *International Transactions on Systems Science and Applications* 4(2).

Thomas, J., and Young, R. 2006. Author in the loop: Using mixed-initiative planning to improve interactive narrative. In *Proceedings of the ICAPS 2006 Workshop on AI Planning for Computer Games and Synthetic Characters*.

Weld, D. 1994. An introduction to least commitment planning. *AI Magazine* 15.

Young, R.; Riedl, M.; Branly, M.; Jhala, A.; Martin, R.; and Saretto, C. 2004. An Architecture for Integrating Plan-Based Behavior Generation with Interactive Game Environments. *Journal of Game Development* 1.