Improving Plan-Based Interactive Narrative Generation

Justus Robertson
Liquid Narrative Group
Department of Computer Science
North Carolina State University
Raleigh, NC 27695

Abstract

I am interested in algorithms that generate well structured interactive narratives that allow human participants to effect change in a story world. Narrative can be structured as a sequence of actions to be carried out by a set of characters in a particular environment. This narrative structure can be generated by AI planning. A process called *accommodative mediation* may be used on a single narrative plan to create a participatory experience that branches based on a human's interactions as a character within the story. However, accommodation is limited by search time and the number of plans the algorithm can possibly reach. In this paper, I discuss an initial modification of accommodative mediation that expands the space of plans the algorithm may search. I also identify several avenues by which the process may continue to be improved.

Introduction

Interactive narrative systems attempt to tell a well-structured story where user actions affect the flow of story events. In order to balance story structure with user agency, an interactive narrative system must incorporate user actions into the plot and ensure that all possible stories are interesting and engaging. For example, the Façade system (Mateas and Stern 2005) maintains a set of dramatic beats which are broken down into pre-authored dialog between the main characters. The user navigates the possible beat and dialog sequences by providing natural language and gesture input.

Another way to structure an interactive experience is to produce a plan that all actors in a virtual environment are expected to follow (Young 1999). Progress has been made to increase the complexity and quality of these computationally generated narrative plans (Riedl and Young 2004; Ware and Young 2011). However, a human user with sufficient autonomy as a character in the story will be able to act out of accordance with the intended story plan. The process of *mediation* (Riedl, Saretto, and Young 2003; Harris 2005; Riedl et al. 2008) merges harmful user actions with a story's existing plot structure in order to create interactivity. Mediation operates by *intervention* or *accommodation*.

Intervention disables a user's exceptional action such that the intended outcome of a behavior never occurs. This approach is less desirable than accommodation, which incorporates an initially unwanted action into the plot by restructuring the narrative's future events. If intervention repeatedly dismisses actions, a user may realize their behavior is limited. Even so, mediation algorithms are forced to intervene when an accommodative solution cannot be found. Mediation cannot accommodate a user's action if it does not have time to exhaust the search space or when no alternate plan exists. My initial work has improved mediation's ability to accommodate exceptional user actions by widening the space of plans accommodative mediation can find. It does this by allowing the algorithm to replan story events that happen before a user's exceptional action.

I would like to continue improving the mediation process and have identified several avenues by which this may happen. First, my current work is built off *reactive mediation* (Riedl, Saretto, and Young 2003), but can be merged with the later *proactive* algorithm (Harris 2005) that attempts to predict user actions using a plan recognition component. Second, much unnecessary complexity arises from the current algorithm's Partial Order Causal Link (POCL) approach to planning. It may be possible to mitigate these complications by converting the algorithm to use a forward-chaining, state-space planner that tracks causal link information, such as the Intentional Fast-Forward Planner (Ware 2012). Finally, the algorithm is able to replan past events while ensuring story consistency by leveraging a simple model of character/player knowledge. I suspect that a robust model will allow the algorithm to remove more steps before replanning, allowing it to search a wider space of plans in turn.

Related Work

The original reactive mediation algorithm (Riedl, Saretto, and Young 2003) was developed for use with the Mimesis system (Young 2001). It reconciles exceptional user actions with the system's original story plan, generated by a narrative planner, in order to create a unified interactive experience. At plan time, the algorithm identifies every causal link in the story plan that could possibly be threatened by some potential action of the user. For each of these link and action pairs, the system generates a list of ways to mediate the user's unwanted behavior. The system either accommodates the user's action into the plot by replanning the story events from the exceptional action onward, or intervenes and prevents the action from happening at all if no accommodative solution is found.

One drawback to reactive mediation is that it only recognizes harmful behavior once a player executes an exceptional action in the game world. Proactive mediation (Harris 2005) attempts to avoid exceptional behavior altogether by incorporating a plan recognition component to predict the player's entire course of action. The system uses a hypothesis of the user's future actions to replan story events or rearrange objects, agents, and properties of the story world in order to foil a user's exceptional plan before it is performed. This algorithm introduces *proactive intervention*, a series of techniques the system can use to disable any of a series of actions that lead to exceptional behavior.

However, there are two problems with this system. First, if the plan recognition component fails to accurately predict the user's behavior, proactive mediation must resort to its reactive counterpart. Second, a system that inserts objects and otherwise alters the state of a story world may introduce noticeable inconsistencies that negatively impact the user's suspension of belief. In order to ensure this does not happen, the current system has very strict rules about what objects it can modify and when they can be modified during proactive intervention. Incorporating my current work into proactive mediation will improve its fail-safe mechanism and ensure story-world consistency from the user's perspective, which will allow it to identify all aspects of the story world it can modify dynamically while ensuring story consistency.

Current Work

My current work allows reactive mediation to search a wider space of plans during accommodation by leveraging a simple model of character knowledge. This model allows the system to identify what the user will have possibly experienced given a POCL plan that represents an interactive narrative. When an exceptional step is identified, the algorithm removes steps from the plan that are determined to be unobserved by the user's character, and replans with the capability of adding consistent events to the plan that occur before the user's exceptional action. In this way, it navigates the user's set of possible worlds during the replanning process, as opposed to a static world determined by the system. This algorithm is called *bidirectional accommodation*.

Data Structures

The algorithm operates on two data structures, a POCL plan and a set of databases that represent the knowledge of story characters. POCL plans are comprised of partially ordered steps, which represent actions to be taken by characters in a story world. These plans are found with *refinement search*, the process of navigating plan-space by correcting flaws in the current plan until none are left. Once found, a partial plan may have many equivalent *total orderings*. A total ordering of a partial plan is any realization of the plan in a valid linear order.

In addition to the POCL plan that represents an interactive narrative, bidirectional accommodation must maintain a database of knowledge information for some set of characters within the story world. This knowledge database is used to prune away steps from the plan before accommodation and control what steps are added during the replanning process. The knowledge annotation process examines every step in the narrative plan and determines whether the step is observed by the user in any possible total ordering. Currently, a step is determined to be observed by the player if the player is at the location of the step in any total ordering of the plan.

Additionally, we check what effects of steps the player will possibly observe in the story. It is possible that the user may observe the ramifications of some action after it has occurred. A player is said to observe an effect of some unobserved step if the player is at the location where the effect exists sometime after the unobserved step occurs in any total ordering of the story. In this case, the database records the effect of the unobserved step in the player's knowledge database.

Due to the partial order nature of POCL plans, we cannot reason about a single series of events. The structure of the plan forces us to consider all possible total orderings. Because of this, we must reason about all possible ways the story may play out. One disadvantage of considering a plurality of orderings is that it constrains what we can consider unobserved by the player to what is unobserved across all total orderings of the plan. Another drawback is that *inconsistent effects*, effects of unobserved actions that may be observed as true or false in different total orderings, force their establishing step to become artificially observed so as not to introduce an inconsistency in certain orderings of the plan.

Policy Generation

The highest level of reactive mediation is an algorithm that takes as input a narrative plan, analyzes the plan for possibly threatened causal links, and creates a policy for how the exceptional actions will be handled. This algorithm has been modified to prepare plans for bidirectional accommodation by first removing all events ordered explicitly between the exceptional action and goal step. Next, a knowledge database corresponding to the user's character is prepared and used to strip away all unobserved steps from the plan. Finally, the algorithm finds the set of flaws introduced into the plan by removing steps, and sends the prepared plan fragment to bidirectional accommodation.

If a valid plan is found, it is included in the algorithm's policy. The algorithm is recursively invocated to find all threatened links in the newly created plan. Once this process accounts for every possible threatened link, the policy is returned.

Bidirectional Accommodation

Bidirectional accommodation is the system's replanning component, able to find new plans that accommodate an exceptional user action identified during policy generation. Bidirectional accommodation is a modification of UCPOP (Penberthy and Weld 1992) that ensures new steps and their effects are *consistent* with the user's knowledge of the story

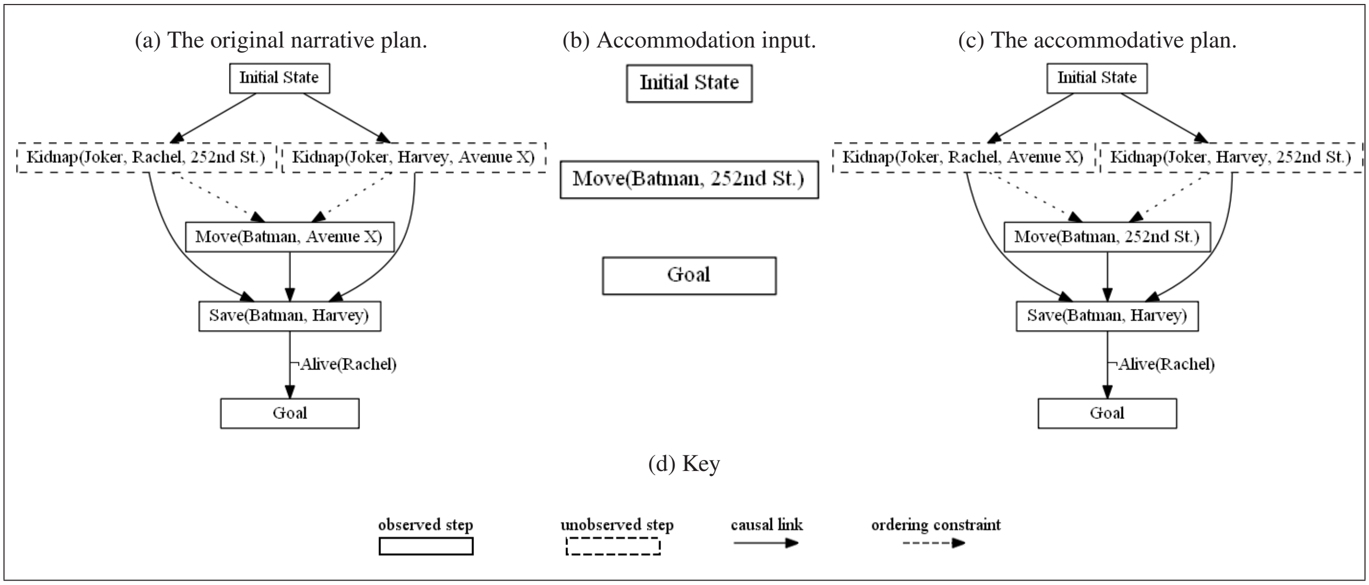

Figure 1: Bidirectional accommodation plans in the Batman domain.

world. A new step is consistent if it is not observed by the player's character in any total ordering of the plan, and an effect is consistent if it does not change any observed state of the world. If successful, bidirectional accommodation returns a story plan that accomplishes the author's objectives, incorporates the user's exceptional action, may include modifications to story events that occurred prior to the enciting incident, and remains consistent with everything the user has observed in the story world.

Example

The third act of *The Dark Knight* (Nolan, 2008) provides an example of how bidirectional accommodation can better provide agency in an interactive experience. While interrogating the film's antagonist, the Joker, Batman learns that both district attorney Harvey Dent and his fiance Rachel Dawes are missing. The Joker tells Batman that he has time to save only one of them. Batman selfishly chooses to pursue Rachel at the expense of Harvey. However, when Batman arrives at Ave. X in pursuit of Rachel, he finds a distraught Harvey. Batman pulls Harvey from the warehouse, but Dent is permanently scarred by Rachel's death. This scene sets into motion the film's final act and is important to the outcome of the story.

If this scene were to play out in an interactive context controlled by reactive mediation where the player takes the role of Batman, they would be given a single choice: move to 252nd St. or Ave. X. If the player chooses 252nd St., they have the opportunity to save Rachel and allow Harvey to die in her place. The system will have no way of accomplishing its goals and will disable the user's action of moving to 252nd St., which removes all player agency from the situation.

However, a model of user knowledge allows the system to recognize events that are outside the user's incomplete knowledge of the story world. These actions are removed from the plan during policy generation and the bidirectional accommodation algorithm is free to find a story that fits Batman's exceptional action of traveling to 252nd St., while ensuring Rachel is the one who perishes in the bomb blast. No matter what the player as Batman chooses to do, they are fated to fulfill the author's goal of rescuing Harvey and ensuring that the final events of the story commence.

Future Work

There are three avenues by which I am interested in expanding plan-based interactive narrative generation using mediation: integrating my work with proactive mediation, migrating the underlying planning algorithm to a forward-chaining state-space planner, and creating a more robust model of observation and character knowledge.

Proactive Mediation

Proactive mediation (Harris 2005) predicts and acts to prevent exceptional user activity before it occurs. By including a plan recognition component, the system can more intelligently intervene to disable harmful user actions by examining the entire sequence of events that lead to a threatened causal link. At its best, proactive intervention changes aspects of the story world to disable exceptional user activity before it begins, instead of reacting to a player's harmful behavior as it happens in the story world. Proactive intervention is comprised of three methods by which the system can prevent exceptional actions: substitution, aversion, and disablement. Of these, I am most interested in aversion.

Aversion in proactive mediation is the process of adding an *inversion* step to the plan that ensures a predicted exceptional action will not occur by changing the state of the world such that one of the action's preconditions cannot be true. The effects of this inversion step are instantaneously added to the story world without simulating any story events that lead up to it taking place. However, this only works if the user has not observed the status of an aversion effect and cannot observe the state of the world changing due to these actions. I suspect that aversion is very similar to the modification I have made in accommodation and that reasoning about a model of user knowledge will allow it to change more aspects of the story world at more times.

State-Space Planning

A limitation of the current accommodation system arises from the partial ordering of POCL plans. While causal links are useful for determining detrimental user actions, they introduce the complexity of many total orderings instead of a single story line. This complexity limits unobserved events or objects to those determined to be unobserved across all possible total orderings, which limits the amount of control accommodation has over the story world. Fortunately, this could be resolved by using a forward-chaining state-space planner, which finds a single total ordering of events. However, state-space planners do not operate on causal links, which is the method by which mediation finds and accommodates exceptional user actions.

In order to take advantage of a state-space planner, causal link information must be accounted for in some way. Fortunately, work has already begun on a forward-chaining state-space planner for generating narrative, the Intentional Fast-Forward planner (Ware 2012). If leveraged for accommodative mediation, a system like IFF could eliminate the inconsistent effects problem and widen the number of steps and effects that mediation can consider unobserved by reducing the set of possible total orderings down to one. This would allow accommodative mediation to have more control over what aspects of the story world it can rearrange. IFF has also been shown to be significantly faster at finding solutions than its POCL counterparts, which should translate to faster accommodation speeds.

Knowledge Model

The current model of character/user knowledge used for accommodative mediation is very simple. It concludes that a character observes a step if the character is at the location at which the step occurs when it occurs. Similarly, if a character is at the location of an effect of an action, the system concludes the character has knowledge of that effect. This may not always be the case for two reasons: not everything at a location is necessarily presented to a character and not everything presented to a character is necessarily observed. This simplistic model is detrimental to the system because it unnecessarily limits the aspects of the story world that accommodative mediation has the opportunity to rearrange.

The first case can be improved by adding more complexity to the existing model of when aspects of the world are presented to a character. For example, in order to know the contents of a chest, a character would not only need to be in the room that contains the chest but open the chest itself and look inside. The second case is harder, because instead of examining the link between what exists in the game world and what is presented to the character, it examines the link between what is presented to the character and what is actually observed. I suspect that not everything presented on a game screen is necessarily observed and internalized by a user. If there is a way to model the relationship between what is presented and what is observed, the system would have more freedom over what aspects of the world it could control.

Conclusion

One way to structure narrative is a series of actions to be carried out by characters in a story world, which can be generated with AI planning. Mediation is the process by which a branching interactive narrative structure can be created from a single story plan. In this paper, I presented a summary of existing mediation algorithms, how I have improved one of these algorithms, and several avenues by which the process of mediation can continue to be improved.

References

Harris, J. 2005. Proactive mediation in plan-based narrative environments. Master's thesis, North Carolina State University.

Mateas, M., and Stern, A. 2005. Structuring content in the façade interactive drama architecture. In *Proceedings of the 3rd Annual Conference on Artificial Intelligence and Interactive Digital Entertainment*, 93–98.

Penberthy, S. J., and Weld, D. S. 1992. Ucpop: A sound, complete, partial order planner for ADL. In Nebel, B.; Rich, C.; and Swartout, W., eds., *KR'92. Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference*. San Mateo, California: Morgan Kaufmann. 103–114.

Riedl, M. O., and Young, R. M. 2004. An intent-driven planner for multi-agent story generation. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 186–193. IEEE Computer Society.

Riedl, M.; Stern, A.; Dini, D.; and Alderman, J. 2008. Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on Systems Science and Applications* 4:23–42.

Riedl, M.; Saretto, C. J.; and Young, R. M. 2003. Managing interaction between users and agents in a multi-agent storytelling environment. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, AAMAS '03, 741–748. New York, NY, USA: ACM.

Ware, S. G., and Young, R. M. 2011. Cpocl: A narrative planner supporting conflict. In *AIIDE*.

Ware, S. G. 2012. The intentional fast-forward narrative planner. In *Eighth Artificial Intelligence and Interactive Digital Entertainment Conference*.

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.

Young, R. M. 2001. An overview of the mimesis architecture: Integrating intelligent narrative control into an existing gaming environment.