

Narrative Mediation as Probabilistic Planning

Justus Robertson¹ and R. Michael Young²

Liquid Narrative Group

[1] Department of Computer Science
NC State University
Raleigh, NC, USA, 27695
jjrobert@ncsu.edu

[2] School of Computing and
The Entertainment Arts and Engineering Program
University of Utah
Salt Lake City, UT, USA 84112
young@eae.utah.edu

Abstract

Narrative mediation is a plan-based process that converts a planning problem into a mediation tree that models an interactive story. Probabilistic planning relaxes the classical planning constraint of deterministic action outcomes by containing domain operators with nondeterministic effects. This paper shows narrative mediation can be viewed as a probabilistic planning problem where the outcomes of user actions are modeled with nondeterministic effects in the domain. The paper gives a process for automatically converting a deterministic mediation planning problem into a probabilistic interactive narrative planning problem.

Introduction

A big picture goal of the field of interactive narrative is to create interesting narrative experiences in an interactive setting where the participant has a strong sense of autonomy and yet the series of events that plays out conforms to a narratively interesting structure. One approach to creating these experiences is a strong story experience manager (Riedl and Bulitko 2013), an intelligent agent that invisibly controls the actions of non-player characters and manipulates virtual world events behind the scenes according to a central data structure that models one or more well-formed stories.

One approach strong story experience managers use to create these data structures is a plan-based process called mediation (Riedl and Young 2006). Mediation begins with a planning domain and problem specified in a modelling language like the Planning Domain Definition Language (PDDL) (McDermott et al. 1998). The problem models an initial state of the world and a desired goal state of the world. The domain models actions that can be performed by story characters to transition the story world from one state to another. Mediation begins by constructing a single plan, or sequence of actions that transforms the initial world state into the goal state, that solves the planning problem. Once it has a single plan, mediation examines the sequence of events and finds alternate actions the player could take during the plan's execution that would break the initial plan's control. For each of these alternate actions, mediation constructs an

alternate plan and recurses until all paths reach a goal state. These plans are connected together in a tree of plans, called a mediation tree, that branches for player action and is used to control interaction during gameplay.

This current mediation process views user actions as something within its control during the planning process. It then uses a second, meta-process to account for user autonomy in order to create a branching mediation tree. An alternate way to view mediation is not as a branching series of deterministic planning problems, but as a single non-deterministic problem where the agent has full control over the outcome of NPC actions but no control over the outcomes of actions taken by the player. This view allows mediation to take the variability introduced by autonomous player actions in its meta tree building process and include it directly in the planning problem by relaxing the deterministic outcome constraint of classical planning. Probabilistic planning problems can be specified in a modelling language like the Probabilistic Planning Domain Definition Language (PPDDL) (Younes and Littman 2004) and solved by a probabilistic planner. The output of a probabilistic planner is a Markov Decision Process (MDP) policy that maps world states to system actions. The rest of this paper outlines related work in narrative mediation and the use of MDPs to solve interactive narrative problems, describes how mediation can be modeled with a probabilistic planning problem, presents a process for automatically converting a classical planning problem to an interactive PPDDL problem, and gives an example of the process in a small domain.

Related Work

Mediation is a plan-based approach to interactive narrative generation. Mediation was first created for the Mimesis (Riedl, Saretto, and Young 2003; Young et al. 2004) system and has been iterated on and incorporated into many systems since then (Riedl et al. 2008; Ramirez, Bulitko, and Spetch 2013; Ramirez and Bulitko 2014; Robertson and Young 2014a; 2014b; 2015b; 2016; 2015a). The common feature of each of these systems is they all build branching story trees by constructing an exemplar narrative by solving a planning problem and then accommodate alternate player actions by constructing branching plans. This paper shows how the accommodation process of mediation systems can be factored into a problem with non-deterministic effects.




Figure 1: A small mediation tree with two story branches.

The output of a probabilistic planner is a Markov Decision Process, a policy that maps states to actions the agent should take during story executing in the virtual environment. MDPs were first used in the context of experience management by the Declarative Optimization-based Drama Management (DODM) system (Nelson et al. 2006a; 2006b), which built on the search-based drama management pioneered by the Oz Project (Bates 1992; Weyhrauch 1997). MDPs serve as the basis for targeted trajectory distribution markov decision processes (TTD-MDPs) (Roberts et al. 2006; Bhat et al. 2007; Cantino, Roberts, and Isbell 2007; Roberts, Cantino, and Isbell Jr. 2007; Roberts et al. 2007). Created to provide replayability in the context of experience management, TTD-MDPs model desired distributions over outcomes and trajectories instead of maximizing expected reward. This paper shows how a mediation problem specified in PDDL can be converted to a probabilistic planning problem specified in PPDDL that accounts for player choices with non-deterministic effects and, when solved by a probabilistic planner, that produces a policy for controlling interaction in an experience management context.

Mediation as Probabilistic Planning

Narrative mediation can be modeled as an MDP where actions taken by NPCs are deterministic and have a 100% transition probability but actions taken by human-controlled interactors are non-deterministic. This paper assumes there is an explicit ordering between action turns for the player and non-player characters in the mediation tree and corresponding MDP. Non-deterministic player actions must be taken by the system and the outcome is split between the possible deterministic actions a player could take in the original domain. The transition probabilities for human-controlled interactors could be predicted with information provided by a model of choice preference (Yu and Riedl 2013), goal recognition (Cardona-Rivera and Young 2015), and/or role assignment (Dominguez et al. 2016). Here we assume players make uniformly random choices among options. This assumption is just a placeholder and can easily be replaced with a more robust model.

Figure 1 shows a small mediation tree. Rows correspond to plans produced by mediator’s planner and branches represent choices the player has between multiple possible actions. S_0 and S_1 are world states, P_0 and P_1 are player actions, N_0 is an NPC action, and G_0 and G_1 are goal states. The story starts at S_0 and the player is given the choice between action P_0 and P_1 . If the player chooses P_0 then the




Figure 2: A small MDP that corresponds to the mediation tree in Figure 1.

system performs N_0 and the story reaches goal state G_0 . If the player chooses P_1 then the story reaches goal state G_1 .

Figure 2 shows a small MDP that corresponds to the mediation tree shown in Figure 1. S_0 and S_1 are still world states and G_0 and G_1 are still goal states. Player actions P_0 and P_1 from the mediation tree have been combined into a single action P_0 with a probabilistic outcome. The first outcome, with probability 50%, corresponds to the mediation tree’s P_0 action and leads to S_1 . The second outcome, also with probability 50%, corresponds to the mediation tree’s P_1 action and leads to G_1 . As mentioned before, the system assumes the player makes uniformly random choices between deterministic actions so two possible action outcomes are each assigned a 50% probability, three would have 33.3%, four 25%, and so on. N_0 is still an NPC action with a deterministic outcome. The MDP also contains an additional NPC action, N_1 , not in the mediation tree because it is not part of any plan created by mediation. A policy for this MDP that corresponds to the mediation tree in Figure 1 is $\pi = \{(S_0, P_0), (S_1, N_0)\}$.

Modelling uncertainty introduced by player choices in plan-based interactive narrative environments as probability distributions over possible action outcomes allows policies to be created by probabilistic planners that correspond to mediation trees. The next section shows how a deterministic PDDL narrative problem can be automatically transformed into a non-deterministic PPDDL problem.

PDDL to PPDDL

In order to encode the non-determinism introduced by game time player actions into a planning problem, we encode fully ground situations in which players can take action directly into probabilistic domain operators. To illustrate this process we use a simple planning problem shown in Figures 3a and 3b. Figure 3a shows the deterministic PDDL planning problem and domain. The player controls a character named Arthur and the computer controls a character named Merlin. Arthur and Merlin are standing at a location named the Woods. A Rock and Excalibur are on the ground at the Woods with Arthur and Merlin. The goal of the planning problem is for Arthur to be holding Excalibur. Characters in this domain can pick up an object if the character and the

Problem

Initial State

Arthur is player
Arthur at Woods
Merlin at Woods
Excalibur at Woods
Rock at Woods

Goal State

Arthur has Excalibur

Domain

take(?taker,?thing,?location)

Precons: ?taker at ?location
 ?thing at ?location
Effects: ?thing not at ?location
 ?taker has ?thing

Domain

Arthur_takes_Excalibur_Woods()

Precons: Rock not at Woods
 Excalibur at Woods
Effects: **Probability 1.0**
 Arthur has Excalibur
 Excalibur not at Woods

Arthur_takes_Both_Woods()

Precons: Rock at Woods
 Excalibur at Woods
Effects: **Probability 0.5**
 Arthur has Rock
 Rock not at Woods
.....
Probability 0.5
 Arthur has Excalibur
 Excalibur not at Woods

Arthur_takes_Rock_Woods()

Precons: Rock at Woods
 Excalibur not at Woods
Effects: **Probability 1.0**
 Arthur has Rock
 Rock not at Woods

take(?taker,?thing,?location)

Precons: ?taker at ?location
 ?taker not Arthur
 ?thing at ?location
Effects: **Probability 1.0**
 ?thing not at ?location
 ?taker has ?thing

(a) PDDL Domain and Problem

(b) PPDDL Domain

Figure 3: A simplified description of a small PDDL domain and problem. A simplified description of a small PPDDL domain that corresponds to the PDDL domain shown in Figure 3a.

object are at the same location. A plan for this domain and problem is the single action (*take Arthur Excalibur Woods*). During gameplay, the player also has the opportunity to perform the action (*take Arthur Rock Woods*).

There are two components of the corresponding probabilistic domain that encodes non-deterministic effects of possible player choices. The first is a copy of the deterministic operators tailored to prevent the player from using the operator and containing a 100% transition probability to model the system's complete control of non-player characters. Figure 3b implements this component with the *take* action containing the added precondition *?taker not Arthur*. The second component is a set of fully ground operators that model possible player choices while playing the game. To generate this set, the set of fully ground player actions possible in the story world must be generated using the deterministic PDDL domain and problem. In the context of the domain and problem presented in Figure 3a, the set of possible fully ground player actions is $\{(take\ Arthur\ Excalibur\ Woods), (take\ Arthur\ Rock\ Woods)\}$.

The next step is to determine intersections between ground player actions. Two actions intersect if there is one or more world states where both sets of preconditions are satisfied. In the case of our example, the two ground player actions intersect when both the Rock and Excalibur are at the Woods. If an intersection takes place, an action should be added to the domain that contains the intersecting preconditions along with probabilistic effects that split the probability across the original effects from the deterministic domain according to a user model, which in our case is simply a uniformly random distribution. The *Arthur_takes_Both_Woods* action is an example of a non-deterministic intersection ac-

tion that splits the effects of the deterministic $\{(take\ Arthur\ Excalibur\ Woods), (take\ Arthur\ Rock\ Woods)\}$ action set.

In addition to intersections, there may be situations where one action is enabled but another is not. In our example, if the Rock has been picked up then the player can only execute the (*take Arthur Excalibur Woods*) action and if Excalibur has been picked up the player can only execute the (*take Arthur Rock Woods*) action. In addition to intersections, the conversion process must also account for situations where possible actions converge. These situations are modeled in the Figure 3b domain with the *Arthur_takes_Excalibur_Woods* and *Arthur_takes_Rock_Woods* actions. Each of these actions have a single outcome with 100% probability because it's the only action the player can perform in the situation.

This process of finding situations in which actions intersect and differentiate is repeated recursively until all choice situations are modeled. The full set of non-deterministic operators has at most one operator representing each combination of single and combination operators. The largest intersecting operator possible is the one that represents all fully ground actions from the deterministic problem.

Conclusion

This paper shows the meta-planning process of mediation can be viewed as a probabilistic planning problem. The paper explains the relationship between narrative mediation and a Markov Decision Process policy that maps states to actions, gives a process that can convert a deterministic mediation problem into an equivalent non-deterministic planning problem, and gives an example of the conversion process.

References

- Bates, J. 1992. Virtual Reality, Art, and Entertainment. *Presence: Teleoperators & Virtual Environments* 1(1):133–138.
- Bhat, S.; Roberts, D. L.; Nelson, M. J.; Isbell, C. L.; and Mateas, M. 2007. A Globally Optimal Algorithm for TTD-MDPs. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, 199. ACM.
- Cantino, A. S.; Roberts, D. L.; and Isbell, C. L. 2007. Autonomous Nondeterministic Tour Guides: Improving Quality of Experience with TTD-MDPs. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, 22. ACM.
- Cardona-Rivera, R. E., and Young, R. M. 2015. Symbolic Plan Recognition in Interactive Narrative Environments. In *The Eight Intelligent Narrative Technologies Workshop at AIIDE*.
- Dominguez, I. X.; Cardona-Rivera, R. E.; Vance, J. K.; and Roberts, D. L. 2016. The Mimesis Effect: The Effect of Roles on Player Choice in Interactive Narrative Role-Playing Games. In *Proceedings of the 34th Annual CHI Conference on Human Factors in Computing Systems*.
- McDermott, D.; Ghallab, M.; Howe, A.; Knoblock, C.; Ram, A.; Veloso, M.; Weld, D.; and Wilkins, D. 1998. *PDDL - The Planning Domain Definition Language*.
- Nelson, M. J.; Mateas, M.; Roberts, D. L.; and Isbell, C. L. 2006a. Declarative Optimization-Based Drama Management in Interactive Fiction. *IEEE Computer Graphics and Applications* 26(3):32–41.
- Nelson, M. J.; Roberts, D. L.; Isbell Jr., C. L.; and Mateas, M. 2006b. Reinforcement Learning for Declarative Optimization-Based Drama Management. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, 775–782. ACM.
- Ramirez, A., and Bulitko, V. 2014. Automated Planning and Player Modelling for Interactive Storytelling. *IEEE Transactions on Computational Intelligence and AI in Games*.
- Ramirez, A.; Bulitko, V.; and Spetch, M. 2013. Evaluating Planning-Based Experience Managers for Agency and Fun in Text-Based Interactive Narrative. In *Artificial Intelligence and Interactive Digital Entertainment*, 65–71.
- Riedl, M., and Bulitko, V. 2013. Interactive Narrative: An Intelligent Systems Approach. *AI Magazine* 34(1):67–77.
- Riedl, M. O., and Young, R. M. 2006. From Linear Story Generation to Branching Story Graphs. *Computer Graphics and Applications* 26(3):23–31.
- Riedl, M. O.; Stern, A.; Dini, D. M.; and Alderman, J. M. 2008. Dynamic Experience Management in Virtual Worlds for Entertainment, Education, and Training. *International Transactions on Systems Science and Applications* 4(2):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 International Conference on Autonomous Agents and Multiagent Systems*, 741–748.
- Roberts, D. L.; Nelson, M. J.; Isbell, C. L.; Mateas, M.; and Littman, M. L. 2006. Targeting Specific Distributions of Trajectories in MDPs. In *Proceedings of the National Conference on Artificial Intelligence*, 1213.
- Roberts, D. L.; Bhat, S.; Clair, K. S.; and Isbell Jr., C. L. 2007. Authorial Idioms for Target Distributions in TTD-MDPs. In *Proceedings of the National Conference on Artificial Intelligence*, 852.
- Roberts, D. L.; Cantino, A. S.; and Isbell Jr., C. L. 2007. Player Autonomy Versus Designer Intent: A Case Study of Interactive Tour Guides. In *AIIDE*, 95–97.
- Robertson, J., and Young, R. M. 2014a. Finding Schrödinger’s Gun. In *Artificial Intelligence and Interactive Digital Entertainment*, 153–159.
- Robertson, J., and Young, R. M. 2014b. Gameplay as On-Line Mediation Search. In *Experimental AI in Games*.
- Robertson, J., and Young, R. M. 2015a. Automated Gameplay Generation from Declarative World Representations. In *Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*.
- Robertson, J., and Young, R. M. 2015b. Interactive Narrative Intervention Alibis through Domain Revision. In *The Eighth Intelligent Narrative Technologies Workshop at the Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*.
- Robertson, J., and Young, R. M. 2016. A Model of Superposed States. In *Proceedings of the 3rd Experimental AI in Games Workshop at the 12th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 65–71.
- Weyhrauch, P. 1997. *Guiding Interactive Drama*. Ph.D. Dissertation, Carnegie Mellon University Pittsburgh, PA. CMU-CS-97-109.
- Younes, H. L. S., and Littman, M. L. 2004. PPDDL1.0: The Language for the Probabilistic Part of IPC-4. In *International Planning Competition at the 14th International Conference on Automated Planning and Scheduling*.
- Young, R. M.; Riedl, M. O.; Branly, M.; Jhala, A.; Martin, R. J.; and Saretto, C. J. 2004. An Architecture for Integrating Plan-Based Behavior Generation with Interactive Game Environments. *Journal of Game Development* 1(1):51–70.
- Yu, H., and Riedl, M. O. 2013. Data-Driven Personalized Drama Management. In *Proceedings of the 9th AAAI Conference on Artificial Intelligence for Interactive Digital Entertainment (AIIDE)*.