Open-World Planning for Story Generation

Mark O. Riedl
Institute for Creative Technologies
University of Southern California
13274 Fiji Way, Marina Del Rey, CA 90292
riedl@ict.usc.edu

R. Michael Young
Department of Computer Science
North Carolina State University
Raleigh, NC 27695
young@csc.ncsu.edu

Abstract

The ability to generate narrative is of importance to computer systems that wish to use story effectively for entertainment, training, or education. One way to generate narrative is to use planning. However, story planners are limited by the fact that they can only operate on the story world provided, which impacts the ability of the planner to find a solution story plan and the quality and structure of the story plan if one is found. We present a planning algorithm for story generation that can non-deterministically make decisions about the description of the initial story world state in a least-commitment fashion.

1 Introduction

The standard approach to incorporating storytelling into a computer system is to script a story at design time. That is, a system's designers determine ahead of time what the story should be and hard-code the story into the system. An alternative approach is to generate stories either dynamically or on a per-session basis using planning.

Young [1999] suggests that planning has many benefits as a model of narrative.

- Plans are comprised of partially ordered steps. If steps are character actions, then a plan is a good model of a story *fabula* – a chronological enumeration of events that occur in the story world.
- Planning algorithms construct plans based on causal dependencies between steps. This ensures that all events are part of causal chains, resulting in coherent story structure.

Story planners such as Universe [Lebowitz, 1985] and IPOCL [Riedl and Young, 2004] require a predefined world description as input. The task of creating the world in which a story is set is already accomplished by the "human author" – the user – and a story planner only has to find a plot that can be performed in the story world. It is possible for the human author to describe a story world initial state that cannot be transformed through character actions into the desired outcome. In this case, the responsibility for a planner's failure lies with the human author. Another possibility is that the story planner can find a solution plan, but the solution plan is strangely structured because the planner must work around limitations inherent in the initial state description.

It is our belief that a story planner that has discretion over the description of the initial story world state will:

- Be able to avoid failure conditions that are attributed to the human author's definition of the initial world state.
- Find better quality solutions because the planner does not need to work around limitations inherent in the human author's description of the initial world state.

2 Initial State Revision

The use of the closed world assumption (CWA) presupposes complete and correct knowledge about the world. Since the universe of discourse that describes any story world is potentially infinite, the practical implementation of the CWA is to enumerate everything about the world that is known to be true so that everything that is not explicitly declared to be true is presumed false. Since story worlds are often fictional, the CWA places a burden on the human author to completely specify the story world or risk limiting the stories that can be constructed by the planner. However, there may be many aspects of the story world that the human author may not have a preference over.

The Initial State Revision (ISR) algorithm partitions knowledge about the initial world state into *three* sets: known true ($\mathcal{T}$), false ($\mathcal{F}$), and undetermined ($\mathcal{U}$). $\mathcal{T}$ is the set of atomic ground sentences that are known to be true. $\mathcal{F}$ is the set of atomic ground sentences that are known to be false or are assumed to be false because their true value is not specified by the human author. $\mathcal{U}$ is the set of atomic (but not necessarily ground) sentences whose truth values are undetermined but not assumed to be false. $\mathcal{U}$ is the set of knowledge about the state of the world that the human author has explicitly designated as having no preference over. Consequently the possible values of $\mathcal{U}$ define *possible worlds* in which the story can be set. Initial state revision reduces the set of possible worlds by committing to values of sentences in the undetermined set as necessary to produce the best possible story plot.

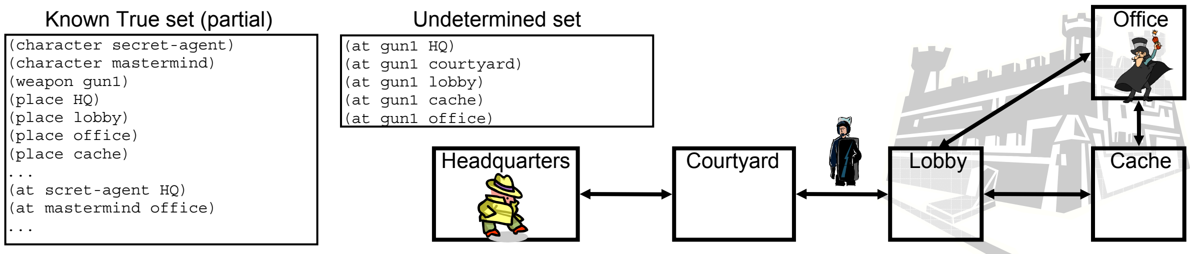

Figure 1: Diagram of the secret agent world.

2.1 Motivating Example

Figure 1 shows an example domain in which a secret agent must assassinate a terrorist mastermind. However, the terrorist mastermind lives in a fortress which is heavily guarded such that no one can enter the fortress while carrying a weapon. The description of the story world disallows any story that involves the secret agent entering the fortress and assassinating the terrorist, for the secret agent will be unable to carry a weapon past the guards. Furthermore, suppose the story planner is constrained such that the terrorist mastermind never leaves the fortress. In this case the story planner can not find a complete story plan for this story world unless the planner is able to revise the initial state of the world. Figure 1 shows the known true and undetermined portions of the initial world state. Properties of the gun have been defined such that its location is undetermined. Through initial state revision, the story planner can declare the gun to have always been inside the fortress, thus making it possible to find a complete plan.

2.2 The Algorithm

ISR planning is an extension of the conventional partial order causal link (POCL) planning algorithm [Penberthy and Weld, 1992]. Plan steps have preconditions and effects. In every complete plan, preconditions of a step must be satisfied by a causal link from the effect of a preceding step or by the initial world state. An open (unsatisfied) condition on a plan step can unify with sentences in any set in the initial step, including the undetermined set, When an open condition p on some plan step unifies with a sentence q in the undetermined set, a causal link is established as in typical POCL planning and, additionally, sentence q becomes determined. If p is not negated, q is determined to be true and moved into set $\mathcal{T}$, otherwise q is determined to be false and moved into set $\mathcal{F}$. By using q to satisfy an open condition, the planner has in effect committed to a set of possible worlds in which the value of q is determined. Note that it is not necessary for the undertermined set to be empty for planning to be complete. That is, the story plan can be in a *set* of possible worlds if the distinguishing details are not integral to the plot.

It is often the case that propositions that describe the state of the world are mutually exclusive, i.e., if one is true then the other cannot be true. For example, an object cannot be in more than one place at a time. The undetermined set $\mathcal{U}$ can contain atomic sentences that should never be simultaneously true in the initial state of the world. ISR utilizes *mutual exclusion sets* (also called *mutex sets*). A mutex set is a set of sentences such that no two sentences can both be true in the initial state. If a sentence in a mutex set is determined to be true, then all other sentences in the mutex set are determined to be false. In the example, the sentences in U are all mutually exclusive to ensure that the gun cannot be in more than one place.

3 Conclusions

Planners are constrained to operate in the world as provided by their users. For an agent operating in a limited domain, it is possible for the world to be completely specified. For infinitely large domains such as a fictional story world, it is impossible to encode every fact about the world. The closed world assumption places a burden on the human author to describe a world that supports story generation. It is possible, however, that despite the human author's intentions, the story world description does not adequately suit storytelling, resulting in planner failure or in solution story plans that are oddly structured to overcome limitations of the world description. The ISR planning algorithm assumes creative control over aspects of the story world description. With control over the initial state of the story world, a story planner can find better story plans because it is not limited by assumptions made by the human author.

Acknowledgements

The work presented here has been supported by the National Science Foundation CAREER Award #0092586.

References

[Lebowitz, 1985] Michael Lebowitz. Story-telling as planning and learning. *Poetics*, 14, 1985.

[Penberthy and Weld, 1992] J. Scott Penberthy and Daniel S. Weld. UCPOP: A sound, complete, partial order planner for ADL. In *Proc. of the 3rd Int. Conf. on Knowledge Representation and Reasoning*, 1992.

[Riedl and Young, 2004] Mark O. Riedl and R. Michael Young. An intent-driven planner for multi-agent story generation. In *Procedings of AAMAS*, 2004.

[Young, 1999] R. Michael Young. Notes on the use of planning structures in the creation of interactive plot. In *Proc. of the 1999 AAAI Fall Symposium on Narrative Intelligence*, 1999.