Rethinking Traditional Planning Assumptions to Facilitate Narrative Generation

Stephen G. Ware and R. Michael Young
Digital Games Research Center
Department of Computer Science
North Carolina State University
Raleigh, NC, 27695

Abstract

STRIPS-style planning has proven to be a helpful methodology for narrative generation, but certain assumptions about the process remain in use which inhibit the creation of interesting stories. The sequence of actions is more important than the initial and goal state of the world, so a narrative planner should first build a plot and then adapt the world to that plot. This is possible by relaxing the closed world assumption to allow revision to the initial and goal states.

Adapting Planning for Story Writing

The mature AI formalism of STRIPS-style planning has proven to be a useful tool for narrative generation. Research groups, especially those headed by Young, Cavazza, Mateas, and Riedl have produced a number of story writing systems based on planners. The STRIPS planning representation is a natural one for story writing because it resembles existing narratology models (Cavazza and Pizzi 2006) and can be easily mapped onto psychological models of story comprehension (Young 1999).

Classical planning was first developed for solving real world problems in which the initial and goal states were known and a sequence of actions to achieve the goal was desired. Most classical domains are assumed to be fully observable, deterministic, static, and discrete (Russell and Norvig 2003). These assumptions are unrealistic for the real world.

However, these assumptions *can* be reasonable for story generation. The author is usually free to take a God's eye view of the story world (fully observable), determine the exact outcomes of actions (deterministic), control the environment (static), and estimate real world concepts to an arbitrary degree of accuracy (discrete). Ironically, classical planning is much better suited to story generation than to its original purpose because of these assumptions. But some traditional assumptions about input and output no longer hold, and those should be critically examined.

Focus on Plot

Much effort has been devoted to adapting planners to produce not the efficient plans they were designed for but *interesting* plans which follow the maxims of good storytelling (especially Porteous and Cavazza 2009; Riedl and Young 2004). While the desired output of the planner—a sequence of actions—is the same, the focus of the problem has changed; the narrative properties of that sequence are now paramount.

The initial and goal states between which those actions materialize impose a lot of constraints on the kinds of plans that can be formed. Even with sufficiently rich guidelines for producing interesting stories, a planner will only consider actions which contribute to the goal and are eventually accessible from the initial state.

In narrative terms, the initial and goal state correspond to the situation of a story world before and after the plot. The focus of narrative generation should be this plot. Since the writer is free to manipulate the world to serve the story, this freedom should be leveraged in story planners as well.

Initial State Revision

Riedl and Young (2005) introduced the Initial State Revision algorithm (ISR), an extension to UCPOP (Penberthy and Weld 1992) which divides the starting state of a plan into three sets: T, the set of true facts; F, the set of false facts plus all facts assumed false via the closed world assumption; and U, a set of facts with undetermined truth value. During plan refinement, ISR moves facts from U into either T or F in order to accommodate the construction of a more interesting plan than would otherwise be generated from a fully specified initial state.

It may seem tempting to do away with the closed world assumption all together—begin with F as an empty set and assume all facts not in T are in U. However, this approach is problematic because nothing prevents the planner from making bad assumptions about the world; i.e. that an item is in two places at once. ISR handles this problem by including with U a set of mutex relations that describe which facts are mutually exclusive. Moving one fact from a mutex group into T requires that all others be moved into F.

This solution is valid though perhaps onerous because it requires the author of the planning problem to provide detailed information about all undetermined facts.

Generalizing ISR - Most General Initial States

Rather than require the problem author to provide a set of mutex relations, we propose that the domain author should provide a Most General Initial State (MGIS) as part of the domain. An MGIS is a partially specified initial state such that all valid initial states for any problem in a domain resolve with the MGIS. In other words, all valid initial states are more specific versions of the MGIS.

An MGIS allows a planner to deduce automatically all mutex relations for a given problem, which can then be given to the ISR algorithm.

With some simple syntactic sugar, an MGIS can be represented very compactly. For example, `(AT ITEM ?X)` can be taken to mean "the item must begin at some location." The planner decides on a convenient binding for `?X` as it searches for a solution.

The use of an MGIS allows one relatively simple logical formula to ensure that ISR never makes bad assumptions for any problem in a given domain. This is especially helpful if the user of a planning system is not the author of the planning domain.

The user can specify as much or as little of the initial state as is desired. If it is fully specified, ISR becomes a classical planner. If no initial state is specified (that is, the MGIS is the initial state), the planner will mold it to fit the plan. Generalized ISR will be most effective if the planner has some knowledge about how to construct interesting plans.

Most General Goal States

A similar notion can be applied to goals states. Rather than specifying, for example, that the villain is dead by the end of the story, the author can simply specify that someone is dead and allow the planner to construct whichever story is deemed most interesting. An MGGS is potentially less useful than an MGIS because authors often desire control over the story ending.

Generalized ISR in UCPOP

Like ISR, this generalized algorithm can be implemented in UCPOP with minimal modification. Because UCPOP binds variables in a least-commitment fashion, variables can be left unbound in the initial state, e.g. `(AT ITEM ?X)`.

An MGIS may include disjunctions, and Weld (1994) notes that disjunctive effects are difficult to handle in UCPOP because they represent a kind of uncertainty. However, in this context the problem can be ignored. Since the planner chooses which fact from a disjunction to make true it need not account for all eventualities. Disjunctive effects can be handled much like disjunctive preconditions.

If the domain also includes an MGGS, this becomes the goal state (preconditions of the last plan step) in a similar manner to how the MGIS becomes the initial state.

Generalized ISR in SAT-Based Planning

Generalized ISR can also be easily implemented in faster SAT-based algorithms. Similar to Kautz and Selman (1996), GraphPlan can be used to explore a plan space and the DPLL algorithm (Davis, Logemann, and Loveland 1962) used to solve the SAT problems derived from the graph.

All facts in U must be included in the GraphPlan initial state because any are potentially true. To ensure that invalid assumptions are not made as a result, the axioms of the MGIS are included in the generated SAT problems.

All logical formulas can be reduced to conjunctive normal form, so this method can be applied to any MGIS. However, the process can be sped up by modifying the DPLL algorithm to handle not only disjunctions but also mutex relations. The change is simple: when one literal in a mutex is set to true, all others must be set to false.

Conclusion

Certain traditional assumptions made in classical planning are helpful to narrative generation because they simplify the problem, but others are detrimental and actually impose unnecessary constraints. Since story authors are free to modify the world, there is no longer a need to fully describe it. Rather, the planner can focus on building an interesting sequence of actions and should be allowed to modify the initial and goal states to fit the story.

This can be accomplished by relaxing the closed world assumption via Initial State Revision. The mutex relations needed for ISR can be automatically deduced by the planner if a domain provides a most general initial state, and ISR planning problems can be solved by classical least-commitment algorithms as well as more modern SAT-based algorithms.

References

Cavazza, M., and Pizzi, D. 2006. Narratology for Interactive Storytelling: A Critical Introduction. In *TIDSE*. 72–83.

Davis, M.; Logemann, G.; and Loveland, D. 1962. A machine program for theorem-proving. *Comm. of the ACM* 5(7):397.

Kautz, H., and Selman, B. 1996. Pushing the envelope: Planning, propositional logic, and stochastic search. In *AAAI*.

Penberthy, J., and Weld, D. 1992. UCPOP: A sound, complete, partial order planner for ADL. In *Knowledge Representation and Reasoning*, 103–114.

Porteous, J., and Cavazza, M. 2009. Controlling Narrative Generation with Planning Trajectories: The Role of Constraints. In *ICIDS*, 234–245.

Riedl, M., and Young, R. 2004. An intent-driven planner for multi-agent story generation. In *AAMAS*, 186–193.

Riedl, M., and Young, R. 2005. Open-world planning for story generation. In *IJCAI*.

Russell, S., and Norvig, P. 2003. *Artificial intelligence: a modern approach*. Pearson.

Weld, D. 1994. An introduction to least commitment planning. *AI magazine* 15(4):27.

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