Collaborative Authoring of Plan-Based Interactive Narrative

James M. Thomas

Department of Computer Science
North Carolina State University
Raleigh, NC 27695 USA
jmthoma5@ncsu.edu

Abstract

My research provides an interface for non-technical authors to collaborate with a planning system to create interactive narrative. I describe a domain metatheory to allow for qualitative elaborations of narrative domains. A graphical user interface that exploits this metatheory is used to specify authorial preferences. These preferences are employed to enhance the qualitative reasoning of the planning system.

Research Problem

"Interactive narratives" are the stories that develop within virtual worlds in which human users interact with one or more computer controlled agents. A persistent challenge in these systems is the narrative paradox: "how to reconcile the needs of the user who is now potentially a participant rather than a spectator with the idea of narrative coherence." (Aylett 2000).

My research builds on an approach for solving this paradox first described as the Mimesis system (Riedl, Saretto, and Young 2003). Mimesis generates plans for actions of story world characters based on hierarchical task decompositions and discrete causal requirements. Although Mimesis simultaneously solves for plot coherence and character believability, the authors acknowledge (Riedl and Young 2004) that a primary limitation is the lack of a search space heuristic that would allow the system to judge the relative "goodness" of one plan over another. In other words, there is no mechanism to ensure that particular narrative qualities such as "suspense", "surprise" or "romance" will be produced in resulting plans.

One might attempt to define a generalized heuristic function in terms of universally accepted narrative ideals, but most planners lack a sufficiently powerful model to make associations between such generalized ideals and the semantics of a specific problem domain and plan space. Even if that bridge were built, there is no consensus view of the ideals that guarantee "good" narrative in the first place. As author Somerset Maugham quipped, "There are three rules for writing the novel. Unfortunately, no one knows what they are".

Research Strategy

My plan is to define heuristic functions for each interactive narrative based on the author's preferences of setting and story. For the system to capture these preferences and report them to the planner, it must have an integrated understanding of the definitions of actions and entities in the problem domain (the setting) and the effects that the constraints on those actions have in defining the topology of the plan space (story experiences). To best create that integrated understanding, my research asks a patient author to remain "in the loop" throughout the plan construction process.

A major challenge of this strategy is how to best conserve and apply the limited time and attention of the author. One of the artifacts of my research strategy is a GUI tool called Bowman that allows the plan author to:

1. Describe the domain.
2. Describe the goals for the story.
3. Describe selection criteria (heuristics).
4. Request possible story plans.
5. View graphs of story plans.
6. Compare story plans.
7. Refine and reiterate.

Bowman is part of the Zócalo suite of planning tools available at NCSU at http://zocalo.csc.ncsu.edu. Like Mimesis, the planning component of Zócalo is based on Longbow, a decompositional (HTN) partial order causal link planner described by Young, Pollack and Moore (1994). To allow the author to "request story plans" as shown above, Bowman passes an XML representation of the planning problem to a planning web-service to generate plans. The Bowman user may specify an arbitrary URL at which the web service resides, or the user may direct plan request to an on-board instance of the planner bundled with Bowman. The planner interface supports several styles of plan requests. The user may request ask for the next complete plan, or for the next N plans, or for as many plans as can be generated by in N seconds. The Bowman user can explore the details of individual plans as well as the entire plan space through scalable vector graphics (SVGs) that can be navigated through mouse over and mouse click. For example, as shown in figure 1, each step in the plan is represented by a collection of rounded rectangles and labeled arrows between steps represent causal and ordering links.

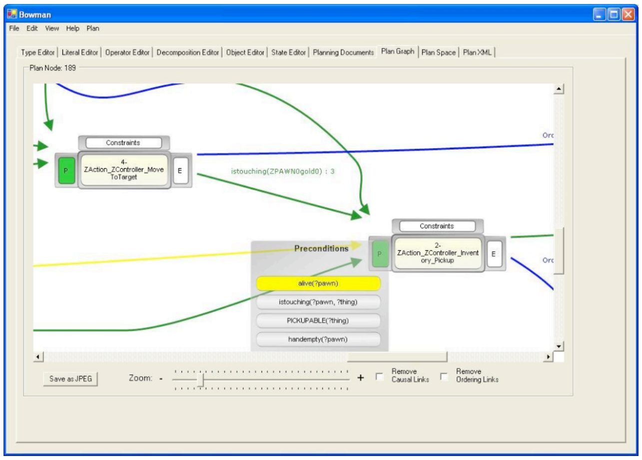

Figure 1: Bowman - Plan Node View

The smaller rectangle on the left side of each step labeled with the letter "P" contains the preconditions of the step. The precondition bubble for a step is filled in with a green colored background if all the preconditions are satisfied for the plan. If the user moves the mouse over the precondition bubble, a semi-transparent window pops up containing each of the preconditions in the plan. As the user moves the mouse down over each precondition, it is highlighted in yellow, as is the causal link which establishes that condition.

Figure 2 contains a Bowman depiction of the plan space as a tree of nodes, where each node is a partial plan.

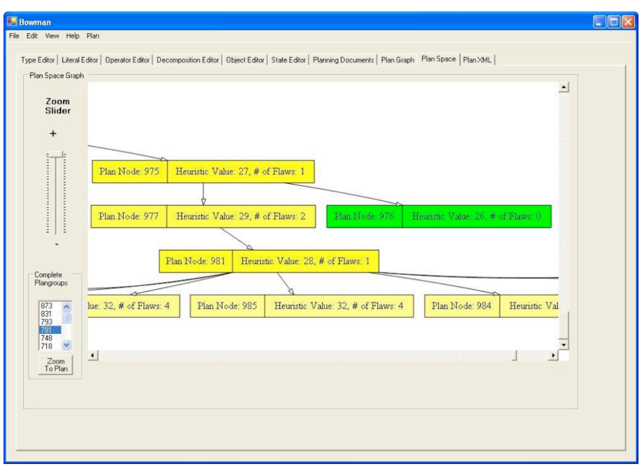

Figure 2: Bowman - Plan Space View

Plan nodes are colored and labeled according to the number of plan flaws they contain. A plan flaw is an open precondition, a threatened causal link, or a flawed decomposition. Plan nodes with zero flaws are shown in green and plans with one or more flaws are shown in progressively lighter shades of yellow. The author can move from the plan space view to view a particular plan node by clicking on it.

Although a modern graphical interface to the planning system certainly affords some efficiency and expressivity advantages, exploration of any reasonably interesting planning problem will quickly exhaust the patience of even the most patient and computer-friendly human author. The core of my research is knowledge representation and elicitation strategies that optimize the use of the human author's creativity. A key requirement is to provide reasonable default strategies for elements of the knowledge representation so that the author need not explicitly express neutral opinions and may focus efforts on those areas for which they are most opinionated. To make this problem more tractable, these strategies are informed by the requirements of planning for interactive narrative.

Interactive narrative domains occupy a promising intermediate level of complexity between the "blocks world" and the real world. Because interactive narrative takes place in a virtual world, its domains are both fully knowable and fully malleable. An advantage for planning research is that these domains may be amended or contracted to suit the requirements of the planning problem. But interactive narrative also introduces special challenges for planning. For example, it is not enough to find a single complete and consistent plan. Authors are interested in understanding how unplanned user actions may affect story goals. This in turn raises issues about the variability of narrative experiences that are possible with each construction and how those possibilities shift as authors make changes.

In fact, the plan author may be responsible not only for the story, but is likely responsible for the domain representation as well. Furthermore, the author may be involved in the creation of the virtual world that is the planning domain. As interactive narrative planning is a component within this larger creative process, there are possibilities and requirements for experimentation and exploration than are not found working with real world domains. This affords researchers new ways to investigate relationships between domains, their representations, planning problems, and the resulting plan spaces. The most fervent hope I hold out for my research is that it leads to new insights into these relationships.

Increasing Domain Knowledge in Planning

Traditional automated planners are not designed specifically to facilitate iterative collaboration with the plan author. Research into collaborative planning methodologies has generally been referred to as advisable or mixed-initiative planning. Advisable planning (Myers 1996) attempts to shape the behavior of the planner by adding additional information to the definition of the planning problem prior to planner invocation. Mixed-initiative planners allow for the iterative and incremental construction of the plan with both the user and the planner capable of proposing or initiating requests to change aspects of the problem or solution. In essence, advisable planning is a special case of mixed-initiative planning, where the initiative is first taken by the plan author, then by the planning system. As such, the success of advisable planning is strongly tied to the knowledge representation it employs to describe the domain. Myers has demonstrated the value of a "domain metatheory" (Myers 2000) that describes the planning domain in terms an author can use to evaluate resulting plans (Myers and Lee 1999). Further research in this area has shown methods for prioritizing the decision choices made available to the plan author to maximize the impact on plan quality (Wolverton 2004). This is a promising method for conserving the limited time and attention of the human author in a mixed-initiative system.

Myers' domain metatheory serves two masters in that it is meant to be intelligible and relevant to human plan authors but also serves as a basis for automated reasoning about plans. My research tool, Bowman, translates abstract domain metatheoretic constructs presented to human authors into a more compact representation for use by the planner. I call this representation the Domain Elaboration Framework, or DEF. The basis of DEF is a STRIPS-style (Fikes and Nilsson 1971) planning domain characterized by objects, conditions and operators. Where the domain metatheory introduced by Myers relies on a description logic of roles, role-fills, features, and measures, DEF uses an alternate grammar of **types**, **dimensions**, **weights**, and **measurements**.

A type is a symbolic name of a node in a global hierarchy of author-defined types with a unique root node named "***anyThing***". Every operator, parameter, and object instance has at least one associated type, and zero or more associated **measurements**. A **measurement** consists of a **dimension** and a **weight**. A **dimension** is a symbolic name selected from a global list of unique author-defined dimensions. A **weight** specifies a relative intensity of the dimension on a normalized interval. Thus, DEF provides a very simple and general elaboration of planning domains at a fine level of granularity. An application employing DEF must provide greater expressive power at the level of the user interface.

Application of DEF to Interactive Narrative

A series of conventions, mostly enforceable through the Bowman GUI, facilitate the application of DEF to the domain of interactive narrative.

Agent Types

A mechanism is needed to distinguish between user-controlled agents and system-controlled agents. System controlled agents are often referred to as bots, or as NPCs - Non Player-Controlled characters. This distinction can be realized through a convention applied to the population of the type hierarchy of DEF. For interactive narrative domains, the type hierarchy can be rooted with "agent" and "non-agent". "Agent" can be further subdivided into "NPC" and "User". Bowman can ensure that all operators contain explicit representations of the types of agent capable of invoking the action.

Mediation Strategies

As described in by Riedl, Saretto, and Young (2003) the planner is responsible for detecting user actions that could threaten the story plan. For each of these *exceptional* actions, the system must determine if changing part of the unexecuted portion of the plan can *accommodate* the action or if an *intervention* is required. An *intervention* requires that the requested action does not execute. Instead an instance of a non-threatening action, called a *failure mode* is substituted for the requested action in real-time.

Bowman can expand the depiction of complete plans to include the application of all available mediation strategies. The plan author can use Bowman to compare these expanded complete plans to see how resilient each is to user action. Authors may be interested in ensuring that the alternative narrative paths dictated by alternative user actions are different or similar based on various qualitative criteria.

Bowman may also help the author in the definition of failure modes. Since a failure mode is simply a list of operators, Bowman can easily highlight the subset of operators from the current library that are good candidates for a particular failure mode, or it may prompt the author to invent suitable operators by describing their characteristics. For example, there may be a need for a failure mode that causes a *shoot(?shooter, ?gun, ?target)* action to fail such that *^alive(?target)* is **not** an effect of its invocation.

Custom Heuristics

Bowman allows the plan author to construct heuristics based on any of the attributes described in DEF to apply relative weights on different flaws or features of a plan. Thus, the author can encode arbitrary narrative preferences and use iterative refinement of the plan space to ensure that optimal levels of "kissing" are in each story, ensure that the possible execution paths have the desired level of conformity or diversity, or simply understand the shortest and longest success paths through a narrative.

Narrative Goal Conflicts

As authors build more narrative goals into their planning problems, it may become more difficult to find complete plans, if these preferences are treated as universally hard constraints. In the narrative domain, it is likely that authors would prefer a sub-optimal plan to having no solution at all. Bowman uses the DEF vocabulary to enable the plan author to specify degrees of "softness" of lower criticality to goals. Methods are being explored to iteratively resolve conflicts between narrative goals and the rules that describe the problem domain. This resolution must be in the form of an 'anytime' algorithm, as the size of the decision space likely exceeds the patience of any human author.

Research Status

Currently, Bowman can be used to define planning domains and planning problems using classical constructs of objects, conditions and operators. The only DEF construct to be realized in Bowman to date is "type". Figure 3 shows how Bowman allows the preconditions of an operator (in this case "LaunchSpaceship") to be edited. Note the use of "type" to color-code the parameters of the literals and the object hierarchy. This aids the author in dragging and dropping properly typed object instances or operator parameters to fill these literal definitions. As shown earlier, Bowman can send these problem definitions to a planner and navigate the resulting plan space.

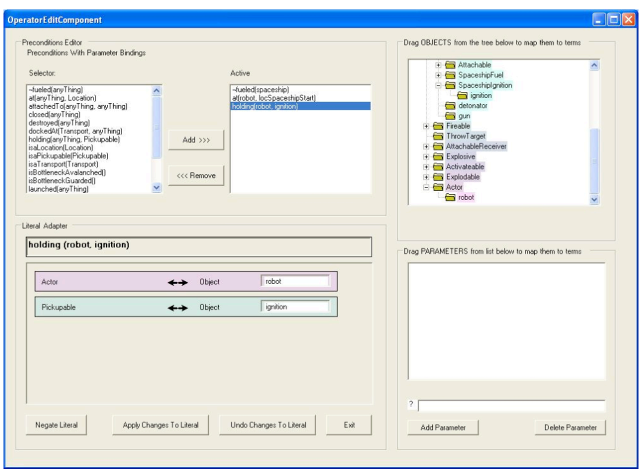

Figure 3: Bowman – Edit Operator Preconditions

Still to be implemented are the remaining DEF constructs, narrative mediation strategies, custom heuristics and higher-level abstractions for managing conflicting hard and soft goals. Finally an evaluation of the expressivity and usefulness of Bowman and DEF must be undertaken. Given the lack of models to which this can be easily compared, its usefulness is likely best gauged through an ablative study to show gauge incremental changes in effectiveness as features are added or removed. I hope to finish this work and defend it in a dissertation within twelve months.

Acknowledgements

This is material is based upon work supported under a National Science Foundation Graduate Research Fellowship and National Science CAREER award 0092586.

References

Aylett, R.S. 2000. Emergent Narrative, Social Immersion And "Storification" *Proceedings, Narrative Interaction for Learning Environments*, Edinburgh.

Fikes, R., and Nilsson, N. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intellligence,* 2, pp. 189-208.

Myers, K. 2000. Domain Metatheories: Enabling User-Centric Planning. In *Proceedings of the AAAI Workshop on Representational Issues for Real-World Planning Systems*, K. Myers and Y. Gil, Eds.

Myers, K., and Lee, T. J. 1999. Generating Qualitatively Different Plans Through Metatheoretic Biases. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence,* AAAI Press.

Myers, K. L. 1996. Advisable Planning Systems. In *Advanced Planning Technology*, A. Tate, Ed. AAAI Press, Menlo Park, CA.

Riedl, M., Saretto, C., and Young, R. M. 2003. Managing Interaction Between Users And Agents In A Multiagent Storytelling Environment. *In Proceedings of the Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-03).*

Riedl, M., and Young, R. M. 2004. An Intent-Driven Planner For Multi-Agent Story Generation. *In Proceedings of the Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-04)*.

Wolverton, M. 2004. Prioritizing Planning Decisions In Real-World Plan Authoring. In *Proceedings of the ICAPS04 Workshop on Connecting Planning Theory with Practice.*

Young, R.M., Pollack, M.E., and Moore, J.D. (1994). Decomposition and causality in partial-order planning. In *Proceedings of the Second International Conference on AI and Planning Systems*, 188-193, Chicago, IL, 1994.