Story Planning as Exploratory Creativity: Techniques for Expanding the Narrative Search Space

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

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

Abstract

The authoring of fictional stories is considered a creative process. The purpose of most story authoring is not to invent a new style or genre of story that will be accepted by the population but to invent a single narrative that is novel enough to be tellable. Computational story generation systems are more limited in the space of narratives than human authors because it is often the case that a story generation system is constrained to operate within a fixed representation of the story world. These limitations can impact whether a story generation system is considered creative or not. In this paper, we describe a story planning system, Fabulist. Fabulist however is constrained by the world model input by the system user. We present two algorithms that enable story planning system such as Fabulist to break outside the bounds of the initial world model in order to search a larger space of narratives.

1 Introduction

Narrative as entertainment, in the form of oral, written, or visual stories, plays a central role in our social and leisure lives. The act of storytelling is commonplace, whether the story is autobiographical – relating something that happened earlier in the day – or fictional. The ability to create and to tell a story is an important and innate skill in humans. There is value then in building computational systems that can create and tell stories. The authoring of fictional stories is considered a creative process. The issue of building story generation systems is twofold: can we build a system that generates stories that are considered "good" and, furthermore, can this computational process of generating stories be considered creative?

Boden [1990] provides a taxonomy of creative systems that distinguishes between *exploratory creativity* and *transformational creativity*. Concepts are locations in *conceptual space* and creativity is the act of identifying new locations within this space. Exploratory creativity is the process of searching an area of conceptual space governed by certain rules. Transformational creativity is the process of transforming the rules and thus identifying a new sub-space. In this framework, media creation such as story authoring is largely exploratory creativity [Turner, 1995; Peinado et al., 2004]. The purpose of most story authoring is not to invent a new style or genre of story that will be accepted by the population but to invent a single narrative that is novel enough to be tellable.

Narrative is the recounting of a sequence of events that has a continuant subject and constitutes a whole [Prince, 1987]. This implies structure to the sequence such as causal relationships between events and a general sense of completeness. A *story* is a narrative that has additional structure that conforms to audience expectations. Examples of structure include dramatic arc and authorial intent (the point of the story). The conventions of story structure significantly constrain the space of narratives that are considered tellable. On top of the issue of the content of the story and how it is structured, there is the issue of narration – the telling of a story to an audience through natural language or cinematic presentation. The work presented here addresses the generation of story content with the assumption that any story content that is generated can be told effectively to an audience. In this work, we consider story generation as a process of finding a sequence of events that can occur in a fictional story world that is considered satisfying to an audience.

The constrained nature of narrative space makes it appealing for search-based story generation. However, computational story generation systems are further constrained by limitations inherent in the representation of the story world in which the story is set. Typically, the representation of the world – what the world looks like, what characters are situated in the world, etc. – is provided as initialization parameters by the system user. We refer to the user of a story generation system as the *human author*.

The remainder of the paper is structured as follows. First, we introduce some related work in computational story generation systems. We then discuss planning as a technology for story generation, how it can be extended to search a larger space of stories with greater potential for finding "good" story plans, and how story plan search relates to Boden's conceptualization of creativity. This discussion raises some issues as to how creative a story planning system can be. The next two sections describe extensions to story planning that enable a story planner to violate the constraints imposed by a human author. The first discusses how a story planner can assume creative control over the description of the story world in which a fictional story is told. The second discusses how a story planner can relax the constraints imposed on it by the human author's given model of the story world dynamics.

2 Related Work

Tale-Spin [Meehan, 1976] is a story generation system that uses an inference engine to determine how individual story world characters will react given the environment and an initial condition. Tale-Spin does not necessarily take into consideration the author's goals. The Universe system [Lebowitz, 1985] uses planning to generate story fragments that do achieve the author's goals. However, Universe emphasizes causal coherence without explicitly addressing other important aspects such as character intentionality.

Case-based reasoning (CBR) techniques have also been applied to story generation. CBR is considered one possible model of computational creativity capable of transformational creativity. Minstrel [Turner, 1994] takes a problem-solving CBR approach, transforming cases from other domains to story situations. Díaz-Agudo et al. [2004] and Peinado et al. [2004] use knowledge-intensive CBR to compare instances of fairy tales based on the grammar of Russian fairy tales developed by Propp [1968].

3 Story Planning

Dehn [1981] argues that a story – an amalgamation of all the events that happen in the story world – must be satisfying or meaningful to the audience and that one way to achieve this is to have the story reflect the intentions of the (hypothetical) author. One way to ensure that a story reflects the intentions of the author is to insist that a specific outcome occurs. Planning is one technology for generating the content of a story that achieves a specific outcome in the form of a goal state.

Young [1999] suggests that planning has many benefits as a model of narrative. The first benefit is that plans are comprised of partially ordered steps. If the plan steps represent actions that are performed by characters in the story world, then a plan makes a good model of a story fabula – the chronological enumeration of events that occur in the story world between the time the story begins and the time the story ends.

The second benefit is that planning algorithms themselves are general problem solvers that attempt to find a sequence of operations that transform an initial world state into a desired goal state. If the initial state of the planning problem represents the state of the story world before the story begins and the goal state of the planning problem is a partial description of the state the story world after the story is complete, then a planning algorithm is model of story authoring that treats authoring as problem solving. The problem is to find a story that can be told in a given story world that results in a given outcome – the goal.

Finally, planners such as UCPOP [Penberthy and Weld, 1992] construct plans based on causal dependencies. Plan steps have preconditions – propositions that must be true in the world for a step to execute successfully – and effects – propositions about the world that are changed by successful completion of the step. Causal dependency planning ensures that all character actions are part of a causal chain of events that lead to the outcome of the story, resulting in a coherent story structure.

Given this argument, any partial order planner can be a story generator given (a) an appropriate world state description, (b) an appropriate library of operations that story world characters/agents can perform, and (c) a method of transforming a plan into written or spoken discourse (e.g. [Callaway and Lester, 2002]) or visual discourse (e.g. [Jhala, 2004]).

Given that a plan shares many of the same properties as a narrative, any plan that is generated by a planner and then narrated can be considered a narrative. Partial order planners such as UCPOP [Penberthy and Weld, 1992] generate plans by searching the space of partial and complete plans. A partial order planner is given a description of the initial world state, a goal (in the case of story planning, an *outcome*), and a library of operators that can be executed by story world characters in order to change the world. Plan operators are discrete actions, represented as a precondition and an effect. The precondition of an action describes all the facts about the world that must be true for an action to be performed. The effect of an action describes how the world is changed by successful execution of the action. Table 1 shows an example of a discrete action schema.

The planning process begins with an empty plan. Planning is the process of identifying and non-deterministically repairing flaws in a candidate plan. Initially, each sentence in the goal statement is an *open condition flaw* that must be satisfied. Each open condition flaw is addressed by non-deterministically selecting an existing action or instantiating a new action from the library of operators. A *causal link* connects two steps s_1 and s_2 via condition e, written $s_1 \xrightarrow{e} s_2$, when s_1 establishes condition e in the world needed by subsequent action s_2 in order for s_2 to execute. When a new action is instantiated in the plan, its preconditions are designated open condition flaws that must also be satisfied. The process iterates by backward chaining until there are no more flaws in the candidate plan. In a complete plan, each precondition of each action is supported by the effects of a previous action or by the initial world state.

Table 1: An action schemata.

```
(define (action shoot)
 :parameters (?attacker ?victim ?weapon ?place)
 :precondition ((character ?attacker) (character ?victim)
                (weapon ?weapon) (location ?place)
                (at ?attacker ?place) (at ?victim ?place)
                (has ?attacker ?weapon)
                (violent ?attacker))
 :effect ((not (alive ?victim))))
```

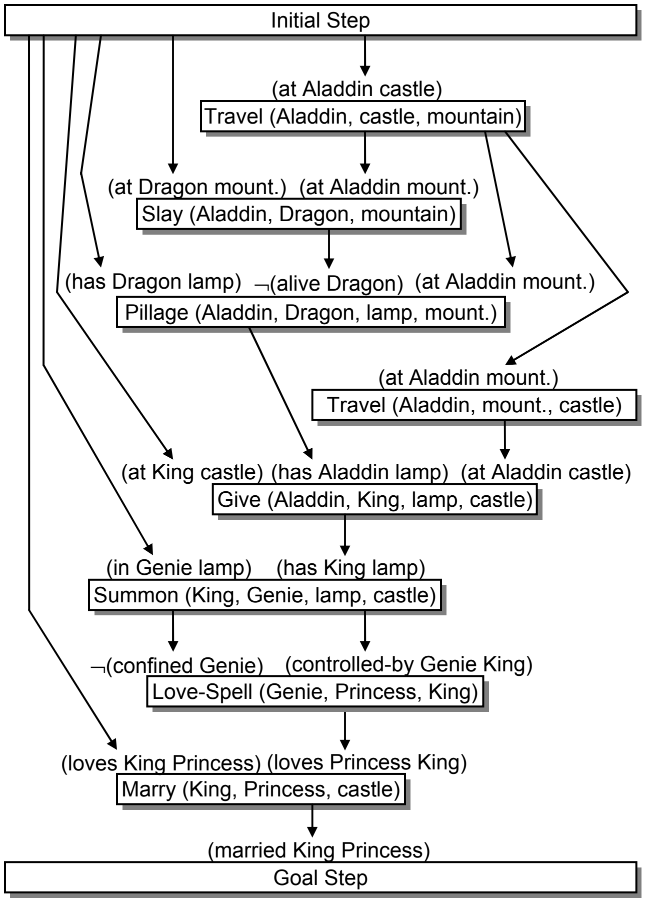

Figure 1: A narrative plan generated by a conventional planner.

The causal links ensure that causal coherence is maintained throughout the plan. That is, each action in the narrative plan enables future actions in the plan to be possible and nothing can happen without apparent cause. Figure 1 shows a narrative plan generated by a planner to "solve" the problem of having a King marry a Princess with the help of Aladdin and a Genie in a magic lamp.

3.1 Story Planning for Character Believability

There is more, however, to story generation than causal coherence of action. Stories also have *character believability*. Character believability refers to the numerous elements that allow a character to achieve the "illusion of life," including, but not limited to, personality and intentionality [Bates, 1994]. Specifically, the actions of believable characters should appear intentional – to be motivated by that character's internal beliefs, desires, and goals. Story generation systems based on conventional planning algorithms have traditionally been unable to generate narratives with character believability [Riedl, 2004]. For example, the narrative plan in Figure 1 does not address the issue of why Aladdin chooses to help the King.

Fabulist [Riedl, 2004] is a story generation system that uses planning technology. Specifically, Fabulist uses the Intent-driven Partial Order Causal Link (IPOCL) planner to generate narratives in which character actions appear causally coherent and intentional [Riedl & Young, 2004; Riedl, 2004]. The IPOCL algorithm addresses the character believability limitation of conventional planning as a technique for story generation by decoupling the characters' intentions from the author's intentions. IPOCL does not assume that the story world characters intend the outcome (goal state) of the story. Instead, IPOCL (1) searches for the intentions that each character might have and (2) motivates through story events why those characters have the intentions that they do.

IPOCL breaks a narrative plan down into sequences of character intentionality – called *intervals of intentionality* – such that all character actions must be part of some sequence. Each sequence achieves an internal goal held by a story world character. An action that is part of an interval of intentionality is considered intentional and all actions must be intentional or the narrative plan is considered flawed. When a newly instantiated action is instantiated into the plan, the planner reasons about why the designated character would perform that action. The action can either (a) achieve some internal character goal or (b) causally support another action that achieves an internal character goal. In the first case, the planner non-deterministically chooses an effect of the action as the goal of a new interval of intentionality. In the second case, the planner attempts to add the action to an existing interval of intentionality for that character.

IPOCL also motivates why a character might have an internal goal and consequently why that character has an interval of intentionality. Each sequence of intentionality must be motivated by the effect of a story-world action. In the previous example, Aladdin might be considered to have the goal that the King possesses the lamp. Why does Aladdin have this goal? IPOCL addresses this flaw by inserting a new *motivating action* in the plan that causes the character to have a goal. In the example, this could be an action in the beginning of the narrative plan in which the King orders Aladdin to acquire the lamp for him. Figure 2 shows a portion of a narrative plan with intentionality represented. The dashed enclosure distinguishes from the rest of the plan (not shown) the Aladdin character's sequence to achieve the goal that the King has the magic lamp. Note that this sequence itself has a precondition – that Aladdin has an intention – which needs to be satisfied.

The extensions to conventional planning enable IPOCL to search a larger space of narrative plans to find candidate solutions that are both causally coherent and character believable with the assumption that those narratives are more likely to be regarded as "good" by an audience. The question, however, is whether IPOCL as a story generator is considered creative? If not, how can IPOCL specifically, and story planners in general, be extended to be considered creative?

3.2 Story Planning and Creativity

Wiggins [2001; 2003] formalizes Boden's description of creativity. The formalization includes the following elements:

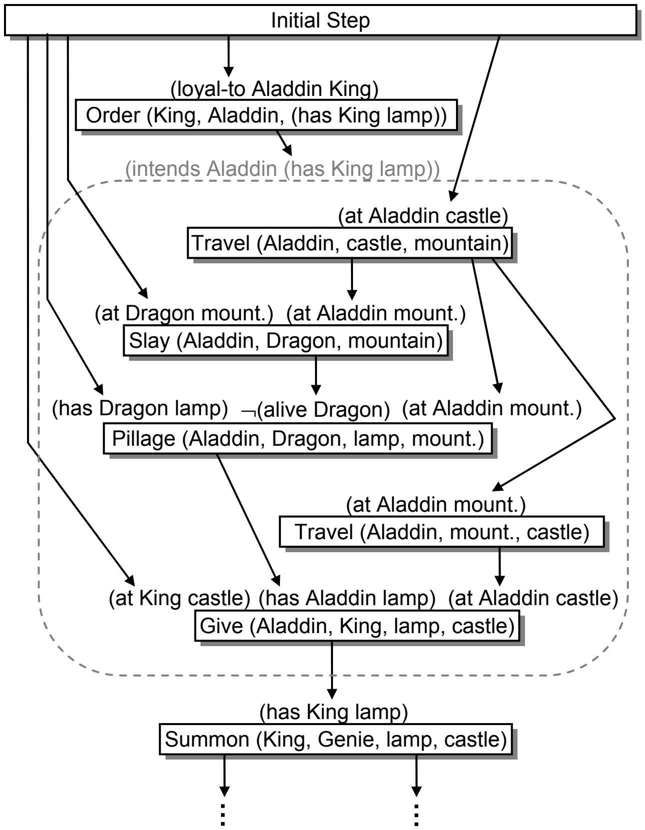

Figure 2: Part of a narrative plan with intentionality.

- $\mathcal{U}$ – The space of all possible concepts relevant to a given domain.
- $\mathcal{L}$ – A language from which to define rules.
- [[.]] – A function generator that maps a set of rules in $\mathcal{L}^*$ to a subset of a conceptual space.
- $\langle\langle ., .\rangle\rangle$ – A search engine for traversing a conceptual space given a set of rules in $\mathcal{L}$.
- $\mathcal{R}$ – A set of constraints in $\mathcal{L}$ that defines a "relevant subset" of $\mathcal{U}$.
- $\mathcal{T}$ – A set of rules in $\mathcal{L}$ that defines a search strategy.
- $\mathcal{E}$ – A set of rules in $\mathcal{L}$ that is used for evaluation of concepts.

Let $\mathcal{U}$ be the space of all narratives, both partial and complete. Let $C = [[\mathcal{R}]](\mathcal{U})$ be a conceptual space that contains the set of "reasonable narratives" for a given domain. $\mathcal{T}$ is a set of rules encoding a search engine that allows the traversal of C such that

$$c_{i+1} = \langle\langle \mathcal{R}, \mathcal{T}\rangle\rangle(c_i). \quad (1)$$

Furthermore if $\perp$ is the empty concept and $^\diamond$ is an iterator function such that $F^\diamond(X) = \bigcup_{n=0}^{\infty} F^n(X)$ then

$$\langle\langle \mathcal{R}, \mathcal{T}\rangle\rangle^\diamond(\{\perp\}) \quad (2)$$

defines a subset of $\mathcal{U}$ that can be traversed by an exploratory creative system. Wiggins (2003) defines an aberration $\mathcal{A}$ as follows:

$$\mathcal{A} = \langle\langle \mathcal{R}, \mathcal{T}\rangle\rangle^\diamond(\{\perp\}) \setminus [[\mathcal{R}]](\mathcal{U}) \quad (3)$$

An aberration is a set of concepts that can be generated by the exploratory creative system but are unexpected solutions, e.g. outside the defined conceptual space. Exploratory creativity can be thought of as the process of finding aberrant solutions that are valuable, that is, the set

$$\mathcal{V}_{\mathcal{A}} = [[\mathcal{E}]](\mathcal{A}). \quad (4)$$

A story generator must traverse many overlapping conceptual spaces including causal coherence, character motivations, themes, and world knowledge (Boden, 1990). Let $C = [[\mathcal{R}]](\mathcal{U})$ be the intersection of these conceptual spaces – a specific space that defines a particular type of story – e.g. a story set in a particular fictional story world. C is therefore the set of all stories that can be told with a particular set of characters in a particular setting. Confounding the issue, story generators operate necessarily within finite computational approximations of story worlds $\mathcal{R}_F \subset \mathcal{R}$ such that $C_F = [[\mathcal{R}_F]](\mathcal{U})$ is a subset of C. In terms of story planning, $\mathcal{T}$ typically encodes a planner such that any narrative $x \in \langle\langle \mathcal{R}_F, \mathcal{T}\rangle\rangle^\diamond(\{\perp\})$ that is not in C_F is automatically rejected whether complete or not. That is, a planner that generates narrative plans can produce novel narratives within a given conceptual space but cannot transgress into new, albeit nearby, conceptual spaces in $\mathcal{U}$. For a system to be creative, it must break the boundaries established by the given initial representation [Haase, 1995].

The initial representation given to a story planner is a model of the story world, consisting of an initial world state description, outcome, and operations that can be performed in the world by characters. The work presented here describes two extensions to the IPOCL narrative planning algorithm that allows IPOCL to violate the constraints of the initial story world representation and generate aberrant solutions. Whether or not the aberrant narratives are valuable is an open question as the evaluation criteria, e.g. $\mathcal{E}$, is not generally known or knowable in the domain of storytelling.

4 Escaping the Constraints of the Initial World Description

Story planners such as IPOCL [Riedl and Young, 2004; Riedl, 2004] and Universe [Lebowitz, 1985] model the story generation process as the search for a plot – a sequence of character actions or events that alter the story world in a way that achieves some goal or some effect on the audience. IPOCL, like most planners, requires a predefined world description as input. The task of creating the world in which a story is set in is already accomplished by an external agency and a story planner only has to find a plot that can be performed in the story world. We define a *human author* as the user of a story generation system that specifies a description of the initial state of the story world and the outcome – a partial description of the state of the story world after the story concludes.

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. It is also possible that the story planner can find a solution plan, but the solution plan is strangely structured because the planner must work around limitations of the initial state description, resulting in a narrative that is rejected by the audience. Wiggins [2003] refers to these cases as *generative uninspiration*. But in these cases, failure can be attributed to the human author.

The technique described in this section – Initial State Revision (ISR) – is a form of open world planning where a portion of the story world is left undetermined by the human author [Riedl 2004; Riedl and Young 2005]. Instead of using sensing actions to determine the true state of the world before instantiating character actions, an initial state revision planner can choose the true state of the world so as to facilitate the story planning process.

4.1 Motivating Example

Suppose the story world provided by the human author contains two principal characters: a secret agent and an international terrorist mastermind. The mastermind lives in a fortress which is heavily guarded. The mastermind is alive in the initial world state and the outcome that the human author desires is a world where the mastermind is no longer alive. Furthermore, suppose the human author specifies the story world such that no lethal weapons are allowed to enter the fortress by the guards. The human author may not have a strong disposition about the exact location of the gun, but is forced to declare every detail of the story world absolutely. Suppose the human author initially places the only gun at the secret agent's starting position, outside the fortress.

The given 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 implements an algorithm such as IPOCL and cannot find a plan in which the mastermind would want to leave the fortress (indeed the mastermind is motivated *not* to leave the fortress). In this case the story planner can find no complete story plan for this story world.

To circumvent problems caused by the way in which the story world was specified, the human author must reason about the impact the description of the world has on the plan search space with regard to the quality and existence of solution plans. The human author would have to have a mental model of the planning algorithm to realize the problem beforehand or to understand the cause of failure afterwards. That is, due to generative uninspiration of the story planner, the human author is forced to debug the story world description. In this case, the human author must correct a determination of a detail she made – the location of the gun – which she did not necessarily have a strong commitment to in the first place.

4.2 Indeterminism in the Story World

The solution to the example problem above is to decouple the story planner's reliance on the initial world state description. Since the universe of discourse that describes any world is potentially infinite, the closed world assumption (CWA) is the assumption that every fact about the world that is not declared true is, by default, not true. However, it is reasonable that there may be certain facts about a story world that a human author does not wish to commit to being true or false. Instead, the planner is given the authority to determine the truth of these facts as it sees fit. Such an indeterminate world state defines a set of possible worlds, each world in the set differing on their truth assignments to the undetermined sentences describing the initial state of the world.

From the perspective of the audience, the story can take place in any one of a universe of possible worlds. Given the audience's uncertainty about the story (which differs from the indeterminate world state representation of the planner, if at all), the audience can consider the actual story world to be one of a set of possible worlds until a fact is revealed in the story that they were once ignorant of. Over time, the set of possible worlds that the audience believes the story is set in converges on the possible world that the author intends, assuming the narration is reliable [Ryan, 1991].

In ISR planning, we define the initial world state as consisting of sentences about the world state that are (a) known true, (b) known false, and (c) undetermined[1].

4.3 Initial State Revision Planning

In partial order plans such as those created by IPOCL [Riedl and Young, 2004; Riedl, 2004] and UCPOP [Penberthy and Weld, 1992], causal links represent the causal relationship between actions in a plan. Each precondition on each action in a narrative plan is satisfied by a causal link from a preceding step. The base case of flaw revision process establishes an unsatisfied precondition on a plan step by extending a causal link from the initial step, s_0, whose effects define the initial state of the world. No step can precede the initial step.

Initial State Revision planning modifies the planning algorithm by allowing certain sentences in the initial step to be indeterminate. When a precondition of an action can be satisfied by an undetermined sentence or the negation of an undetermined sentence in the initial step, the initial world state is revised to make that undetermined sentence known-true or known-false. Thus, the planner makes a commitment to the description of the initial story world state above and beyond what the human author considered. When the initial state is revised, the set of possible worlds that the story is told in is reduced because one source of indeterminism is removed. The initial story world state is considered to have always been defined as such. This does not introduce inconsistencies because no other step has referenced this particular fact in the initial step. If an inconsistency arises somewhere down the line, the planner backtracks and considers a different alternative way of resolving the flaw.

[1] Under the modified CWA, the known-false set of sentences is defined as (universe of discourse) \ (known-true set ∪ undetermined set).

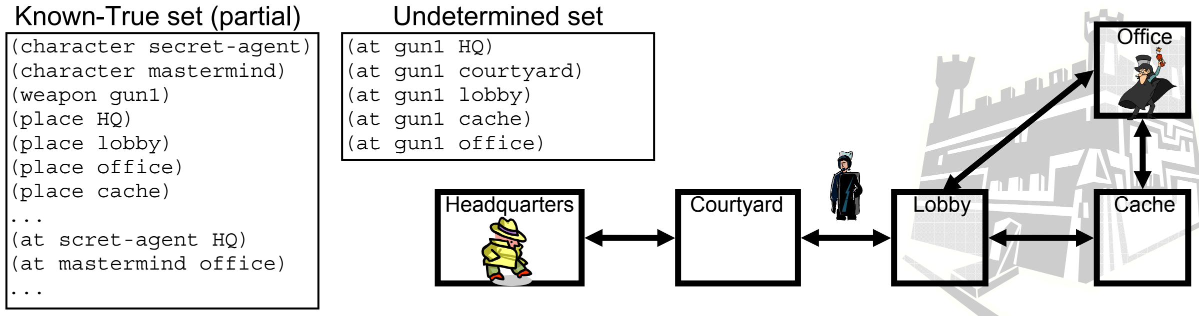

Figure 3: Diagram of the secret agent world.

While initial state revision allows the planner to determine the initial world state that best suites the process of story authoring, it also introduces the potential of logical and semantic inconsistency to the world description. It is often the case that sentences in the initial world state are mutually exclusive, meaning that if one is true than the other cannot be true. For example, an object cannot be in more than one place at a time. So if `(at object place`$_1$`)` is true, then `(at object place`$_2$`)` and `(at object place`$_3$`)` and so on cannot be true. Mutual exclusivity is not an issue in conventional planners because mutual exclusivity is implicitly enforced by the plan operators. While this is also true in ISR planning, the undetermined set can contain atomic sentences that should never be simultaneously true in the initial state of the world. To prevent this situation from arising, 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 simultaneously determined to be false.

4.4 Example of ISR Planning

To illustrate the ISR planning algorithm, consider the story world with the secret agent and the international terrorist mastermind. A diagram of the world is shown in Figure 3. The secret agent starts at the `headquarters` and the mastermind starts at the `office`. The secret agent is prohibited from traversing from the `courtyard` to the `lobby` while holding a weapon. The human author has specified that a weapon, `gun1`, is known to exist but the location of the gun is specified as undetermined.

The outcome of the story is specified by the human author to be a world state in which the mastermind is not alive. The following is a trace of a single path through the plan search space. The planner non-deterministically satisfies the goal condition by having the secret agent shoot the mastermind with `gun1`. In order to shoot the gun, the secret agent must have the gun. The planner considers all the places the secret agent could pick up the gun. Suppose the planner non-deterministically chooses that the secret agent pick the weapon up in the cache location[2]. However, the location of the gun is undetermined. Consequently, an initial state revision occurs, determining that `gun1` must have been in the cache to begin with. The sentences in the undetermined set are mutually exclusive, so when `(at gun1 cache)` is made true, all other sentences are made false to ensure that the gun cannot be in more than one place at a time. From this point, the planner has no problem determining a sequence of actions where the secret agent enters the fortress, proceeds to the cache, picks up the gun, and assassinates the terrorist mastermind.

[2] Any of the three locations inside the fortress would have been acceptable.

5 Relaxing the Constraints of the Story World Model

Another way to escape the conceptual space defined by the human author as inputs into a story planner is to relax the constraints imposed by the library of operations that story world characters can perform. The library of operators constrains the way in which the story world can be transformed from one state to another through actions performed by the characters. Specifically, the preconditions of an operator constrain when and how that operator can be used.

Conceptually there are two types of preconditions. Some preconditions constrain the non-changing attributes of operands. Table 1 shows the schema for an action in which one character shoots another with a weapon. The terms `?attacker` and `?victim` are variables that are bound to objects in the story world when the operation is instantiated in the narrative plan. In this case, the preconditions require that these variables be bound to story world characters. The precondition `(violent ?attacker)` constrains the world such that only characters that have the `violent` attribute associated with them can perform the `Shoot` action. Other preconditions constrain the dynamic state of the world. Preconditions such as `(at ?attacker ?place)` and `(at ?victim ?place)` specify that the attacker and victim must be simultaneously in the same location. This class of preconditions models the physics of the world and should be expected to correspond with the audience's understanding of how the world works.

Instead of using preconditions to specify the attributes of operands, we provide these attributes as *recommendations*. Recommendations are identical to preconditions except the planner can choose whether to satisfy the condition or not [Riedl, 2004]. When a recommendation is considered, the planner divides the search space into $n + 1$ branches where there are *n* ways of satisfying the condition plus one branch in which the condition is left unsatisfied. When the planner chooses to ignore a recommendation, then the planner is intentionally using an object – possibly a character – in a way that is contrary to the model held by the human author and presumably the same as the model of the world held by the audience.

5.1 Acting "Out of Character"

We have used this particular relaxation of constraints to implement a character personality model for the IPOCL story planner. Perception of character personality is an important part of how the audience actively evaluates and predicts future outcomes of a story [Gerrig, 1993]. We assume that the audience makes attributions of broad personality traits to story world characters based on behaviors they observe. The human author can designate personality traits – descriptive terms about consistent behavioral tendencies such as "violent" – to each story world character. The story planner will then attempt to choose actions for those characters that are consistent with those attributions, within reason.

Character personality attributes are encoded in the initial world state and operators recommend which personality traits the character operands should have to perform those actions in the story world. The advantage of using recommendations instead of preconditions for character personality attributes is that it allows the planner to consider narrative plans in which characters act "out of character." When an action uses preconditions to encode character traits, it reduces the space of narrative plans that can be found to be only those in which characters perform actions that are sanctioned according to their traits [Rizzo et al., 1999]. Narratives in which characters act "out of character" due to duress are considered more dramatic. An example of an operator with character trait recommendations is shown in Table 2.

It is possible that the story planner will choose to ignore all recommendations. We provide heuristic functions that favor plans in which more recommendations are satisfied because these are less likely to violate the expectations of the audience. With this heuristic rule, the planner will choose plans in which characters act "out of character" only when absolutely necessary to continue the plot. This comes from the notion that people act consistently except when circumstances require otherwise.

Declarations of character personality traits by the human author, of course, limit the flexibility of the story planner. Fortunately, character trait recommendations can be combined with ISR planning. Character personality traits can be left undetermined in the initial story world state; the fact that action schemata recommend traits instead of require (via preconditions) traits does not fundamentally change the way planning occurs. Recommending undetermined conditions, however, does result in an expanded search space. In the case that a newly instantiated action recommends a particular character trait and that character trait is undetermined, the planner non-deterministically chooses between three possibilities:

- The recommendation is satisfied and the character trait is determined through initial state revision to be known true.
- The recommendation is not satisfied and the character trait is determined through initial state revision to be known false.
- The recommendation is not satisfied and the character trait is determined through initial state revision to be known true.

The first two possibilities are straight-forward. The third is required for algorithm completeness due to the possibility that actions in the story world can conceivably cause a character's traits to change. The planner must consider the possibility that even though the recommendation of character trait matches the initial world description (after initial state revision), something can happen in the interval that causes the character to express the opposite trait and consequently act "out of character."

Table 2: An action schemata that recommends character traits.

```
(define (action shoot)
 :parameters (?attacker ?victim ?weapon ?place)
 :precondition ((character ?attacker) (character ?victim)
                (weapon ?weapon) (location ?place)
                (at ?attacker ?place) (at ?victim ?place)
                (has ?attacker ?weapon))
 :recommendation ((violent ?attacker))
 :effect ((not (alive ?victim))))
```

5.2 Further Constraint Relaxation

A story planner can use recommendations to further relax the constraints of preconditions that model the physical model of the story world. For example, the `Shoot` action schema can *recommend* that the attacker and the victim be collocated. In this case, the story planner is explicitly choosing to ignore the causal relationships between actions that make a narrative causally coherent. Since causal coherence is an important part of the understandability of a narrative and therefore likely to be expected (or even required) by the audience, violating these preconditions will lead to aberrations that are likely to not be considered "good." However, It has been noted that audience members are, to a certain degree, willing to infer what happens during omissions in a story to "fill in the gaps" [Gerrig, 1993]. Again, heuristic functions should be provided to discourage the planner's abuse of recommendations so that recommendations are only ignored when the effect on the audience will be favorable or when the story planner cannot find a solution without violating a constraint.

6 Conclusions

The ability of a story planner to generate satisfactory stories is largely dependent on its ability to search the large space of narratives for those that contain the structural features that most closely conform to audience expectations. IPOCL expands the search space of conventional planners to account for two critical components of story, causal coherence and character believability, which are considered important aspects of meaningful and understandable stories. However, to be considered creative, a story generation system cannot be completely constrained by the story world representation given by the human author.

We present two additional enhancements to story planning technology that lessens the implicit constraints of the initial world representation. Initial State Revision planning enables a story planner to assume creative control over the definition of the initial story world state. Recommendation planning enables a story planner to assume control over the representation of actions.

There are many types of creativity [Boden, 1990]. Story generation seems to be a process that more often than not falls within the bounds of *exploratory creativity* [Turner, 1995; Peinado et al., 2004]. Story generation systems – especially story planners – are constrained by a provided story world representation and this appears to have an impact on whether story generators are considered creative or not. As such the creative aspect of generating fictional stories comes from expanding conceptual space to defy the expectations of the intended audience.

Acknowledgements

This work has been supported by National Science Foundation CAREER award 0092586. Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

References

[Bates, 1994] Joseph Bates. The role of emotion in believable agents. *Communications of the ACM*, 37(7), 1994.

[Boden, 1990] Margaret Boden. *The Creative Mind: Myths and Mechanisms*. New York: Basic Books, 1990.

[Callaway and Lester, 2002] Charles Callaway and James Lester. Narrative prose generation. *Artificial Intelligence*, 139(2), 2002.

[Dehn, 1981] Natalie Dehn. Story generation after Tale-Spin. In *Proc. of the 7th International Joint Conference on Artificial Intelligence*, 1981.

[Díaz-Agudo et al., 2004] Belén Díaz-Agudo, Pablo, Gervás, and Federico Peinado. A case based reasoning approach to story plot generation. In *Proc. of the 7th European Conf. on Case Based Reasoning*, 2004.

[Gerrig, 1993] Richard Gerrig. *Experiencing Narrative Worlds: On the Psychological Activities of Reading*. New Haven: Yale University Press, 1993.

[Haase, 1995] Kenneth Haase. Too many ideas, just one word: A review of Margaret Boden's *The Creative Mind: Myths and Mechanisms*. *Artificial Intelligence*, 79, 1995.

[Jhala, 2004] Arnav Jhala. *An Intelligent Cinematic Camera Planning System for Dynamic Narratives*. Masters Thesis, North Carolina State University, 2004.

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

[Meehan, 1976] James Meehan. *The Metanovel: Writing Stories by Computer*. Ph.D. Dissertation, Yale University, 1976.

[Peinado et al., 2004] Federico Peinado, Miguel Ancochea, and Pablo Gervás. Automated control of interactions in virtual spaces: A useful task for exploratory creativity. In *Proc. of the 1st Joint Workshop on Computational Creativity*, 2004.

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

[Prince, 1987] Gerald Prince. *A Dictionary of Narratology*. Lincoln: University of Nebraska Press, 1987.

[Propp, 1968] Vladimir Propp. *Morphology of the Folktale*. Austin: University of Texas Press, 1968.

[Riedl, 2004] Mark O. Riedl. *Narrative Planning: Balancing Plot and Character*. Ph.D. Dissertation, North Carolina State University, 2004.

[Riedl and Young, 2004] Mark O. Riedl and R. Michael Young. An intent-driven planner for multi-agent story generation. In *Proc. of the 3rd Int. Joint Conf. on Autonomous Agents and Multi Agent Systems*, 2004.

[Riedl and Young, 2005] Mark O. Riedl and R. Michael Young. Open-world planning for story generation. In *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence*, 2005.

[Rizzo et al., 1999] Paolo Rizzo, Manuela Veloso, Maria Miceli, and Amedeo Cesta. Goal-based personalities and social behaviors in believable agents. *Applied Artificial Intelligence*, 13, 1999.

[Ryan, 1991] Marie-Laure Ryan. *Possible Worlds, Artificial Intelligence, and Narrative Theory*. Bloomington: Indiana University Press, 1991.

[Turner, 1994] Scott Turner. *The Creative Process: A Computer Model of Storytelling*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1994.

[Turner, 1995] Scott Turner. Margaret Boden, *The Creative Mind*. *Artificial Intelligence*, 79, 1995.

[Wiggins, 2001] Geraint Wiggins. Towards a more precise characterisation of creativity in AI. In *Proc. of the ICCBR 2001 Workshop on Creative Systems*, 2001.

[Wiggins, 2003] Geriant Wiggins. Categorising creative systems. In *Proc. of the IJCAI 2003 Workshop on Creative Systems*, 2003.

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