R. Michael Young, Stephen Ware, Brad Cassell, Justus Robertson

Plans and Planning in Narrative Generation: A Review of Plan-Based Approaches to the Generation of Story, Discourse and Interactivity in Narratives

1 Introduction

The last ten years have seen a significant increase in computationally relevant research seeking to build models of narrative and its use. These efforts have focused in and/or drawn from a range of disciplines, including narrative theory (Bal, 1997; Chatman, 1980; Ryan, 1991; Herman, 2007), game studies (Jenkins, 2004), computational linguistics (Elson & McKeown, 2010), cognitive psychology (Zwaan & Radvansky, 1998; Magliano, Dijkstra, & Zwaan, 1996; Gerrig, 1993), film studies (Branigan, 2005), multi-agent architectures (Osborn, 2002), and others.

Many of these research efforts have been informed by a focus on the development of an explicit *model* of narrative and its function. Computational approaches from artificial intelligence (AI) are particularly well-suited to such modeling tasks, as they typically involve precise definitions of aspects of some domain of discourse and well-defined algorithms for reasoning over those definitions. In the case of narrative modeling, there is a natural fit with AI techniques. AI approaches often concern themselves with representing and reasoning about some real world domain of discourse – a *microworld* where inferences must be made in order to draw conclusions about some higher order property of the world or to explain, predict, control or communicate about the microworld's dynamic state. In this regard, the fictional worlds created by storytellers and the ways that we communicate about them suggest promising and immediate analogs for application of existing AI methods.

One of the most immediate analogs between AI research and narrative models lies in the area of reasoning about actions and plans. The goals and plans that characters form and act upon within a story are the primary elements of the story's plot. At first glance, story plans have many of the same features as knowledge representations developed by AI researchers to characterize the plans formed by industrial robots operating to assemble automobile parts on a factory floor or by autonomous vehicles traversing unknown physical landscapes. As we will discuss below, planning representations have offered significant promise in modeling plot structure. Equally as significantly, however, is their ability to be used by intelligent algorithms in the *automatic* creation of plot lines. Just as AI planning systems can produce new plans to achieve an agent's goals in the face of a unanticipated execution context, so too may planning systems work to produce the plans of a collection of characters as they scheme to obtain, thwart, overcome or succeed.

Further planning-related work by AI researchers and computational linguists bears relevance to the telling of stories. Significant work in natural language processing has developed computational models of language use as planned, intentional action. Building on work in the philosophy of language (Searle, 1969b; Grice, 1989), these approaches leverage planning systems to reason about the function of individual speech acts in the achievement of a speaker's communicative goals. Building on these models and informed by the rhetorical function of communicative actions, they automatically generate both the content of a discourse as well as its organization. Increasingly, these methods are finding application in the generation of narrative communication, ranging both across the production of narrative text as well as the control of an automated virtual camera used to create a cinematic telling of a narrative within a virtual world or 3D game environment.

There are, of course, contributions to the processes involved in the *understanding* of narrative made by models of plans and activity within story worlds. Some of the key representational issues are discussed by Mueller's (2014) article in this special issue. While these contributions are significant, we focus here on those approaches that have demonstrated progress on the generation side. Our intent is not to suggest that the two tasks are independent. Quite the opposite, in fact. As we discuss below, some approaches to the effective generation of narratives benefit from explicit models of narrative understanding, exploiting a model of a reader's comprehension process to tailor a narrative's content and organization.

While there are many fortuitous similarities between AI plans and the kinds of plans we see play out in narrative, there are also significant differences. These differences are often due to features of narrative that are at the core of storytelling. As a result, they can't be easily overlooked or worked around. While many advances have been made in narrative generation by using AI planning algorithms, a significant amount of research is still required to create rich, expressive plan structures that capture the goal-directed behavior of a story and its telling.

1.1 A Little Context

1.1.1 The Planning Problem

While AI research has seen 50 or more years of effort in the sub-field of planning, with quite divergent approaches and problem specifications, at the core of the planning task is the construction of a sequence of actions in pursuit of a set of goals. Typically, a planning problem is composed of a specification of an *initial state*, that is, a description of a given world as it currently is configured, a specification of a *goal state*, that is, a partial description of the way that we want the world to be, and a *library* of descriptions for all the types of actions that can be performed within the world. To solve a planning problem, a planning algorithm constructs a plan, that is, a sequence of actions instantiated from its library such that, when the first action in the plan is executed in the initial state and each subsequent action is executed in correct order, the resulting world state will be consistent with the goal state description.

Typically, the initial and goal states are described in some restricted form of first-order logic. Action descriptions in the plan library very often follow a general pattern used in early planning systems (e.g., Fikes & Nilsson, 1971) where each action is defined in terms of a set of *preconditions* and a set of *effects*. An action's preconditions are a set of logical terms completely specifying the conditions that must hold in the world immediately prior to the action's execution in order for the action to succeed. An action's effects are a set of logical terms specifying all the ways that the state of the world is changed as a result of the successful execution of the action.

Various modifications and extensions to this base representation have been developed (e.g., actions with probabilistic effects, actions with temporal duration, hierarchically structured actions). While a full review of planning itself is beyond the scope of this article,[1] in the sections that follow we describe just those extensions that have been used in the context of narrative generation.

1.1.2 Practical Concerns

The planning problem, as it is described above and in the planning research community, is known to be P-SPACE complete[2] (Bylander, 1991). The high computational complexity of planning limits its practical applications, especially when much of a machine's processor is already devoted to other tasks such as graphics processing. Before describing planning-based approaches to narrative generation, it is important to justify its use relative to less computationally expensive approaches such as scripting and story grammars.

Planners manipulate atomic units of narrative meaning which are minimally constrained by their context. A planner is free to add an action to a plan for many different reasons (M.O. Riedl & Young, 2010), and actions can easily be reused for multiple purposes in a story (Ryan, 1991; Tomaszewski, 2011). Template-based approaches (e.g. scripting) and grammar-based approaches tend to add events to a story for a single purpose, making it difficult to reason about other purposes they might serve.

Planning is theoretically attractive for the same reason it is practically limited— because it provides a vast space of stories to explore. One way to measure the descriptive power of a computational model is to consider how discriminatory it is within the space of stories. Scripting and grammar-based approaches often have a high degree of human storytelling knowledge encoded into their rules to ensure that every possible branch of the space meets certain standards of narrativity. The intelligence in these system generally lies in the designer's cleverly-defined space. Planning-based approaches navigate a much larger space, which is full of non-narrative plans, and select valid solutions based

1 See, however, the recent proceedings of the International Conference on Autonomous Planning and Scheduling and planning-related papers in the proceedings from the annual conferences from the Association for the Advancement of Artificial Intelligence and the International Joint Conference on Artificial Intelligence for recent work in this area.

2 P-SPACE includes all decision problems which can be solved by a Turing machine in a polynomial amount of space, and is believed to contain NP (Arora & Barak, 2009).

on models of human reasoning. In a narrative planning system, more of the intelligence is encoded in the model than in the search space.

There are a number of ways to compromise between generality and speed in planning. Domain-specific heuristics (Kautz, 1998) can significantly speed up planning in a single domain. Hierarchical Task Networks (Erol, Hendler, & Nau, 1994) and Decompositional Planning (Young & Moore, 1994) allow the designer to encode groups of commonly used actions into script-like sequences which can be automatically expanded by the planner to save computation. Also, current research in fast planning search heuristics is making traditional planning more viable for online systems. Planning has seen initial use in limited contexts for story generation in several commercial video games, including F.E.A.R. (Orkin, 2006) and Killzone 2 (Champandard, Verweij, & Straatman, 2009), as well a numerous fully-realized academic interactive narrative virtual environments (Cavazza, Charles, & Mead, 2002a; Pizzi, Charles, Lugrin, & Cavazza, 2007; Porteous, Cavazza, & Charles, 2010a; Thomas & Young, 2010; Ware, Young, Stith, & Wright, 2014).

1.1.4 Organizing Plan-Based Narrative Elements Into Story, Discourse and Interactivity

Any narrative is a complex collection of inter-related components with interdependent function and meaning. The task of developing a set of computational models for the elements of narrative is a daunting one. Where does one start? How does one begin to factor the elements of narrative into distinct units?

Fortunately for computer scientists and others approaching this task, narrative theorists have given these exact questions considerable thought.

While there is disagreement about specific details among narrative theorists, they broadly divide a narrative into two constituent parts: a *story* and a *discourse*. In this article, we will adopt just this division, put forward by Chatman (Chatman, 1980) and others, and use it as an categorizational filter for the rest of this paper, grouping relevant work into one category or the other.

Here, we use the term *story* to denote everything that is present in the world of a narrative: locations, objects, characters, their personalities, mental attitudes and relationships and all actions and events that transpire within the story world. The term *discourse* denotes the telling of the narrative, that is, the medium-specific communicative elements that an author uses to convey the story to a reader, viewer or player. Discourse includes, for example, the lexical choices, sentences and paragraphs used by an author, the background music or camera edits used by a film director and the voiceover narration used in a cutscene by a game writer.

While not all approaches to plan-based modeling of narrative explicitly follow a conceptual division between story and discourse, we find it useful in this article to characterize the relevant work relative to the production of plans that address story world dynamics (that is, story plans) or plans that address authorial communicative goals (that is, plans for narrative discourse). In addition, we add here at third category: *interactiv-*

ity. This category is not addressed by traditional narratologists, although sociolinguists focused on narrative (e.g., De Fina

& Georgakopoulou, 2012) do consider the complex relationships between author and reader. We include this category here because significant work has been done to use plan-based methods to generate narratives for playable media such computer games.

2 Story

As described above, story is the representation of all the people, places, events, and things in a narrative. Traditionally, story has been viewed as the starting point in the narrative pipeline, so it is not surprising that many initial efforts in narrative generation focused specifically on properties of story. Planning technology is especially applicable for modeling story because it provides a formal, generative paradigm for representing a sequence of actions and the dynamics of a story world. In this section we group plan-based models of story into three broad categories based on the key problem they were designed to address: maintaining story coherence, balancing character and author goals, and representing conflict.

2.1 Maintaining Coherence

Young (Young, 1999) first identified parallels between the causal and temporal data structures of certain types of planning algorithms and the representation of story discussed by narratologists (e.g. Bal, 1997).

Specifically, Young discussed the role that hierarchical plans and partial order causal link (or POCL) plans can play in the representation of narrative structure. Hierarchical plans represent abstract actions and the collection of more primitive actions, arranged in a hierarchy, that serve as sub-plans to achieve their goals. POCL-style plan representations explicitly mark the causal and temporal dependencies between actions in a plan that lead to the accomplishment of a set of goals. These explicit structures are used by plan generators to ensure that they produce well-formed plans – ones guaranteed to achieve their goals or to conform to conventional means of performing subplans – and they can also be used to model important narratological properties like the use of genre-stereotypic plot fragments (in the case of hierarchical plans) or to ensure the causal coherence of a story (Christian & Young, 2004).

When we say that a narrative is coherent, we mean to say that it does not violate the audience's expectations about how the world will change based on the actions that take place. Psychologists have indicated that causality and reasoning about change within a story world is a critical part of story understanding. For example, Trabasso and Sperry (1985) found that events in a causal chain that lead to the outcome of the story are more readily recalled by readers than are events that are not on a causal chain to the story's end.

Because of their relatively direct representation of hierarchical and causal narrative structures, POCL and hierarchical planners have been the foundation for several models of story discussed below, including Lebowitz's (1985) UNIVERSE , Riedl and Young's (2005) IPOCL , and Ware and Young's (2013) CPOCL.

Several researchers have addressed issues of coherence in addition to causal and temporal consistency. Niehaus, Li, and Riedl (2011) discuss strategies for avoiding narrative dead-ends in planning. They define a dead end as an event that does not causally contribute to any future events and might thus be seen as unnecessary or unimportant. Li and Riedl (2010) describe an offline story planning algorithm which attempts to eliminate dead ends by making them part of some causal chain or by removing them. Tomaszewski's (2008) Marlinspike system chooses the next scene of an interactive drama by selecting the scene that will best reincorporate past events. This choice is motivated by the Aristotelian principle called *unity of action* (Halliwell, 1987), but can also be seen as avoiding dead ends in the story by reusing the state changes produces by earlier events.

Porteous et al. (2010a) focus not only on the coherence of individual events but also on the coherence of the story trajectory as a whole. They present a method that leverages planning landmarks (Porteous, Sebastia, & Hoffmann, 2001) – states which must be true at some time during the plan – to ensure that the story arc follows certain genre constraints of rising action, climax, and falling action. Pizzi et al. (2007) designed a plan-based narrative generation system which is notable for representing not only the dynamics of the story world but also the emotional states of the characters. Thus, they ensure not only the causal consistency described by Trabasso et al. but also consistent emotional trajectories for their characters.

2.2 Character vs. Author Goals

One of the most frequently discussed problems in story generation is the tension between the goals of the characters and the goals of the author. Graesser, Singer, and Trabasso (1994) demonstrated that an audience comprehends a story more easily when it perceives that characters are taking actions motivated by their character goals. However, these goals may be at odds with the constraints the author has imposed on the story for various rhetorical purposes.

TALE-SPIN (Meehan, 1977) was the first narrative generation system to be driven by character goals. It generated simple fables in which anthropomorphic animals followed pre-programmed plans in order to fulfill their needs. The effects of each character's actions changed the world state and possibly created new character goals that would then be solved in turn. The role of authorial intent in the construction of the story was largely ignored in TALE-SPIN. To address the role of authorial intent in the construction of stories, Dehn (1981) created the AUTHOR system which made the author's constraints a focus of the story construction process. Whereas TALE-SPIN created stories by simulating characters in a story world, AUTHOR began with a set of actions that an author provided as input, then created a story plan around those actions to justify them *post hoc*. TALE-SPIN and AUTHOR suffered corresponding limitations – for TALE-SPIN not all actions taken by characters contributed to the story, and for AUTHOR, characters sometimes acted without clear motivations.

Lebowitz's (1985) UNIVERSE system made the first attempt to reconcile these two approaches. UNIVERSE constructs episodic melodramas such as sitcoms and soap operas using hierarchical planning. High-level actions corresponding to short scenes in an episode are added to the story plan to satisfy the author's goals for that episode. However, the atomic actions which compose those high- level actions are driven by character goals. The resulting stories achieve some balance between character and author goals, but the system is limited by the extensive amount of hand-authorship required. Each scene has to be annotated with which author goals it could satisfy, and each scene decomposition has to be predefined by a human author in order to ensure coherence.

M. O. Riedl and Young (2010) sought to balance character and author goals while retaining the flexibility of a planner that reasons directly about atomic action sequences. Their Intentional Partial Order Causal Link (or IPOCL) planning framework is an extension of classical planning. When defining an IPOCL planning problem, the author must define which characters (if any) are responsible for taking an action. For example, the act of opening a door must be intended by the character pulling the door knob. During planning, the algorithm tracks a frame of commitment for each character goal. A frame of commitment is a subsequences of actions devoted to a specific character goal all of which must be taken by the character who holds that goal. By reasoning about frames of commitment, IPOCL is able to treat character goals and author goals as the same type of constraint during planning. This also allows it to reuse story actions to satisfy both kinds of goals. In short, an valid IPOCL plan is one which achieves the author's goals by only taking actions which are clearly motivated for the characters who take them. One important limitation of the IPOCL framework is that it considers the goal of the planning problem to be the author's goal, so while the author can control the final state of the story world it is difficult to impose other kinds of constraints, such as how certain kinds of goals should be achieved.

2.3 Conflict

Narratologists identify conflict as an essential element of narrative (Brooks & Warren, 1943; Ryan, Herman, & Jahn, 2005; Abbott, 2008; Egri, 2009). Conflict motivates characters to take action (Egri, 2009), structures the discourse (Ryan et al., 2005), and causes the audience to form expectations about the story's outcome (Gerrig, 1993; Abbott, 2008). Numerous computational narrative researchers (Meehan, 1977; Sgouros, 1999; Szilas, 2003; Barber & Kudenko, 2008;

M. O. Riedl & Young, 2010) have also discussed the importance of conflict to the narrative generation process. However, because most planning techniques are not specifically designed for representing narratives, they are designed to ensure that plans are free of any conflicts that might prevent the plan for achieving the goal. This is a highly desirable property when a planning system is autonomously controlling a rover on a distant planet. But narratives without conflict lack tellability. Carbonell (1981) described an early knowledge-based system for writing stories in which the agents intentionally thwarted one another's plans, however this system relies on pre-defined scripts similar to TALE-SPIN's and so does not have the flexibility to reason about the construction of novel sequences of atomic actions. Smith and Witten (1987) used an adversarial planning algorithm similar to mini-max game tree search (von Neumann, 1928) for generating stories with conflict. The antagonist is modeled as a player engaged in a zero-sum game with the protagonist so as to thwart his or her plans during the story. However, this model does not consider the antagonist's motivations and thus can be described as only satisfying the author's goals to create conflict while neglecting an explanation of the antagonist's intentions.

Szilas's (2003) IDtension system generates stories by simulating what is termed *narrative physics* in addition to the traditional physics of the story world. This systems tracks each character's moral principles and uses those morals when choosing an individual character's next action. However, the overall system itself attempts to generate stories in which characters are motivated to violate their moral principles. This internal conflict over one's morals is the driving force behind the drama IDtension seeks to create.

At least three systems have leveraged the structural properties of plans in order to represent conflict. In POCL-style planning representations, a *causal link* is a data structure used to mark a specific causal relationship between two steps in a plan. A causal link between step s_1 and step s_2 in a plan indicates that s_1 establishes through one of its effects a condition in the world required by a precondition of s_2. During the construction of a plan, if ever a third step s_3 can be ordered between two causally linked steps s_1 and s_2, where s_3 undoes the condition in the world established by s_1 for s_2, the causal link between the two steps is said to be threatened and the plan invalid. Gratch and Marsella developed two systems, Émile (Gratch, 2000) and EMA (Marsella & Gratch, 2009), which compared the plans of multiple agents to discover when a step in one agent's plan threatens a causal link in the plan of another agent. These threats were used to create appropriate affective responses for the the agents who were in conflict. Similarly, Ware et al. (2013) leveraged threatened causal links as a representation of conflict when creating the Conflict Partial Order Causal Link (or CPOCL) framework. CPOCL is an extension of IPOCL that allows certain steps in a plan to be marked as nonexecuted. A non-executed step in a frame of commitment represents a step that an agent intended to take but failed to take due to some conflict. Previously, when a character in the IPOCL framework held some goal, they must have either not pursue the goal or successfully achieved it. The addition of non-executed actions allows for the representation of plans which fail or partially succeed due to causal conflicts with the plans of other characters.

3 Discourse

As we've discussed above, narrative discourse includes those elements of the narrative involved in the *telling* of the story. Increasingly, research efforts in narrative generation have targeted not only story generation but also the creation of the content and organization of the communicative elements of a narrative. Planning-based methods for narrative discourse generation often follow key aspects of work done in the area of natural language generation, where a view of language as planned, communicative action (Searle, 1969a) has lead to success in the generation of multi-sentential text in genres other than narrative (e.g., the work of Moore & Paris, 1993, applied to explanation generation).

In our language-as-action view, elements of a narrative discourse serve as actions in a communicative plan, acting in service of a set of communicative goals. Unlike the plans carried out by characters in a story, narrative discourse plans have effects that obtain outside of the story world. Specifically, these communicitive actions in a narrative are designed to affect the mental state of the reader or viewer. Each discourse element, as it is read or viewed, contributes to the reader or viewer's comprehension process and facilitates the construction of a mental model of the story as it unfolds. Not only is this view consistent with linguistic models of speech act theory, but also with several psychological models of narrative comprehension (e.g., Zwaan & Radvansky, 1998; van den Broek, Pugles-Lorch, & Thurlow, 1996; Graesser et al., 1994).

3.1 Leveraging Plan Representations to Represent and Manipulate A Reader's Mental Model of a Narrative

To be effective, the story and discourse of a narrative must work together to create specific mental configurations in readers or viewers as they experience the narrative's progression. In the work we describe below, system designs have assumed that a story line (also represented as a plan) has been created by exogenous components; the task of the narrative discourse generation systems is to work with the fixed story line and a set of communicative resources (e.g., operators that reorder the telling of the story, elide elements of the story or select medium-specific means to convey the story) to produce a discourse that achieves a targeted effect on a reader's comprehension process.

Cheong and Young (to appear, 2008) developed a system that selects a subset of a story's actions to include in its discourse designed to create suspense around a specific point in the story. The system, named *Suspenser*, takes as input a plan data structure that represents the character actions in the story's plot. As a first step in generating suspense, *Suspenser* characterizes each step in the input plan relative to a measure of the step's importance in the plot (Trabasso & Sperry, 1985). Based on these ratings, it creates an initial candidate discourse that includes a minimal number of the most critical steps. *Suspenser* then simulates the process a reader will use when reading that discourse content to gauge the reader's level of suspense. This process uses a planning algorithm as a proxy for the reader's interpretation process: given the skeleton of a story plan present in the candidate discourse, how might a reader fill in the missing plan content of the plot? Motivated by psychological research on the nature of plot-based suspense (Vorderer, 1996), *Suspenser* ranks a candidate discourse higher in suspense when its planning algorithm can find relatively few successful plans to achieve the story protagonist's goals.

To select a discourse with high levels of suspense, *Suspenser* iterates over the possible contents of the discourse, searching for skeleton story plans that lead its reader model to interpretations where the protagonist seems unlikely to achieve his or her goals.

The *Prevoyant* system developed by Bae and Young (2008) generates a discourse for a narrative that includes foreshadowing and flashback designed to invoke the feeling of surprise in a reader or viewer. The system takes as input a plan data structure characterizing a story's plot. Processing is handled by two core components: a generator and an evaluator. The generator and evaluator work in a generate-and-test iterative approach, first generating candidate discourses for a story plan and then testing to see if it will evoke surprise.

The generator creates a candidate flashback by finding the set of steps in the input plan that directly causally support the goals in the goal state. Causal chains that support these steps from the initial state are then identified by tracing the step's preconditions to the effects of the actions that establish them and repeating this until the chains lead to the initial steps of the plan. Those causal chains that do not affect other steps in the plan are identified as candidates for the content of flashbacks, since their causal independence would allow their presentation to the reader to be deferred without making the story plan appear incoherent or unsound. These causal chains can then be ordered after the steps that support the goal to create a flashback. Once one of these causal chains is chosen, the evaluator tests if it will result in surprise. This is accomplished by trying to construct a plan that does not include the causal chain yet still supports the step that is connected to the goal state. If such a plan can be found, *Prevoyant* marks the candidate discourse as lacking suspense and considers candidates that manipulate other causal chains.

Prevoyant's generator constructs foreshadowing in one of two ways. First, it can present at the beginning of the discourse an object or character that is in one of the flashback causal chains. Second, it can present one of the events in the flashback causal chains at the beginning of the discourse but elide important elements from its telling.

Both *Prevoyant* and *Suspenser* are motivated by Gerrig's view of readers as problem-solvers, where readers work to solve the plot (or planning) related problems faced by a story's protagonist. In these two systems, this motivation translates into the use of planning *algorithms* as proxies for problem-solving processes. Ongoing work by Cardona-Rivera and his collaborators (Cardona-Rivera, Cassell, Ware, & Young, 2012) attempts to build explicit parallels between cognitive models of story events and the relations between them and planning data structures and the relations that hold between steps in a plan.

Their model, called Indexter, is a mapping from the cognitive model of narrative comprehension called the event index situation model (EISM) developed by Zwaan and Radvansky (1998) in to a knowledge representation used for characterizing plans and plan reasoning. Psychologists like Zwann and Radvansky posit that readers construct mental models of an unfolding story as they read, composed of collections of individual events. Each event characterizes a situation, and situations are connected to one another based on shared parameters called indexes. In the EISM, each situation is tagged with five indexes: time, space, protagonist, causality, and intentionality. Two events are connected along a particular index if they share the same value for that index (e.g, two actions that occur in the same location are linked along the space index).

Indexter maps situations and their index values into plan steps and the arguments bound to each step's execution. Space and protagonist indexes correspond to location and character arguments for each action in a plan. The time index corresponds to the timing information that places a plan step in temporal order in its plan's execution. The causality index corresponds to the causal annotations made in POCL style plans to indicate when and how a previous step establishes a given step's precondition(s). And intentionality corresponds to the role that a character's actions play in achieving his or her goals (e.g., the explicit annotations of intentionality in the IPOCL plan representation described above).

Just as the EISM model makes predictions about the salience of past events based on a reader's percpetion of a newly occuring one, so Indexter may be used to make predictions about plan elements' salience as a reader or viewer experiences a narrative generated using planning models augmented with its annotations. In this way, a plan generator might search for story plans whose content and ordering prompt certain salience effects and targeted moments in the reader's experience of the story.

3.2 Planning-Based Visual Discourse Generation

The systems described in the previous section reason primarily about the propositional content of a narrative (what elements from the story plan should appear in the narrative) and the narrative's organization (in what order should specific elements of the narrative be presented) in order to achieve effects on the mental state of the reader like suspense, surprise or salience. Plans are at the core of their representation, but their use of plan generation algorithms is, in some sense, secondary to the process of generating the discourse structure.

Two examples of work that uses planning algorithms to create the narrative discourse from whole cloth both focus on the creation of cinematic, visual discourse. Specifically, these approaches generate plans that are intended to drive a virtual camera operating in a 3D game environment. The camera is used to film a story, driven by a separate plan data structure, playing out within the game world.

Darshak (Jhala & Young, 2010) is a plan based discourse generation system that constructs cinematics based on specific cinematic communicative goals. Darshak makes use of a hierarchical planning model (Young, Pollack, & Moore, 1994) in which primitive actions correspond to individual types of camera shots (e.g., apex shots, close up shots). Hierarchical groupings of actions correspond to cinematic *idioms* (Arijon, 1991), patterns or templates of shots and shot sequences developed by Hollywood cinematographers and readily understood by film viewers.

The preconditions and effects of the action descriptions in Darshak express conditions regarding the beliefs of the viewer about the underlying story world. For example, when a primitive shot films the execution of an action in the story performed by a character, the viewer comes to know that the action occurred and that the character was the agent of the action's execution. Darshak's primitive action descriptions also contain temporal constraints that link the timing of their execution to the times of execution of story actions that they film. For instance, an apex shot used to film one duelist firing a pistol at another is temporally constrained so that it begins filming five seconds prior to the start of the gunshot action and ends at or after the time when the gunshot action ends.

Ember (Cassell & Young, 2013) is system, currently being developed, that augments the methods used by Darshak to produce cinematics containing shots that explain character deliberation. The story plans that Ember is used to film contain conventional actions as well as actions that mark the deliberation of a character towards his or her own plans and goals. Because these actions essentially occur in the character's mind, they may be more difficult for readers to correctly identify and reason about. As Ember constructs its narrative discourse plans, it requires that information used in a character's deliberative efforts are salient to a viewer during the shots that film the deliberation itself. As a result, Ember may create cinematic sequences with shots that film conditions in the world that have previously been filmed in order to re-introduce some story world condition into focus in the viewer's memory. This re-filming of propositions from the story world is similar to the generation of informationally redundant utterances in natural language generation (Walker, 1996).

4 Interactivity

As has been shown in the research described in Section 2 above, one significant capability of planning systems when modeling narrative is their ability to automatically create a wide range of novel story plotlines within a given story world. While many of the research systems described here target the creation of non-interactive text-based or cinematic narratives, the potential to automatically create story lines also offers a major benefit to creators of computer games and virtual worlds. In these applications, story content creation comes at a premium due to the high cost of authoring the content and of designing a user's interaction with it.

For interactive systems like games, however, the use of planning to generate story lines presents a particular challenge. Planning systems for narrative generation produce whole sequences of more primitive actions for plot; they make a particular commitment to the creation of a given sequence, and the narrative experience of that plotline is dependent, in part, on the structure of that sequence. When such a story plan is used in an interactive environment, there is an immediate conflict. On the one hand, a player's perception of the coherence of the story line – its beginning, middle and end – are dependent upon the plan's overall structure. On the other hand, a player's perception of his or her agency in the game world is dependent on the sense that he or she can act in the game world to make substantive changes to the unfolding story. In interactive narrative systems, this issue is termed the *narrative paradox* (Louchart & Aylett, 2003), the conflict between the system's ability to present a well-structured, interesting series of events and the freedom of interaction it offers users.

This conflict exists because the amount of story content needed for an interactive narrative grows by the number of meaningful choices the participant can make. If branching narrative is viewed as a story graph (M. O. Riedl & Young, 2006) whose vertices are story content and whose edges are user actions, meaningful user choices are edges that transition to unique vertices. A branching narrative that offers a large number of unique choices will have an equally large number of vertices, which creates a high authorial burden. Luckily, generative AI techniques such as planning have the potential to create highly branching story graphs with little effort.

There are two popular approaches to generating branching story (Mateas & Stern, 2003b): believable autonomous agents and well-structured plot generation. In this section we examine plan-based approaches to each of these methods of interactive narrative generation.

4.1 Planning Character Behavior

One method of generating an interactive narrative experience that branches on user behavior is to define autonomous character agents who plan and act independently to accomplish goals. The *Teatrix* story creation system (Paiva, Machado, & Prada, 2001) is an early example of this method. Teatrix is a system that allows children between the ages of 7 and 9 to create fairy tales by interacting as and with a number of fictional characters such as Little Red Riding Hood and the Big Bad Wolf. Each system-controlled character is assigned a story role and acts according to general, predefined goals associated with its role. The characters are able to perceive changes in the world, update an internal world model, update their goals, and use an automatic planning algorithm to find sequences of actions that will achieve their desired world state.

Another early system that plans character behavior is Cavazza, Charles, and Mead's work on interactive sitcoms (Cavazza, Charles, & Mead, 2002b). The system uses Hierarchical Task Networks (HTNs), a form of hierarchically structured plan representations, to generate the possible behaviors of characters within the sitcom genre. A HTN can be viewed as a network of high level tasks or goals that decompose into sequences of actions to be taken that achieve those goals. Story emerges in the system as the two characters act to accomplish their goals in a 3D game environment. In this system, the user participates as an observer of the unfolding story who has the ability to move around the virtual stage and intervene in the world by helping or hindering the characters as they execute their individual plans. Actions available to the user include the capability to change the state of objects in the story world, for instance, moving objects from their appropriate places on the set. The actors react to changes in the environment made by the user or other characters when those changes conflict with their current plan. They do this by retracting the related planning decisions made by the HTN planner and re-generating new plans by finding new sub-plans that fulfill their current goal in the new story world environment.

Façade (Mateas & Stern, 2003a) is an interactive drama that experiments with character planning in the context of well-defined plot. The system allows its user to act as the dinner guest of a young married couple named Trip and Grace, whose relationship deteriorates as the evening evolves. The user interacts with the story by moving around the couple's apartment and speaking with the two virtual characters via natural language text input. Trip and Grace operate using a reactive planning language descended from Hap (Loyall & Bates, 1991), called ABL. ABL allows the characters to react to the player within the context of well-defined plot events, called beats essentially small, pre-scripted plan fragments. Façade progresses through a series of beats which each contain a number of goals for the characters to carry out through ABL actions. The user influences the progression of the story based on their interaction in the game world as they attempt to reconcile or disrupt the married couple's relationship.

FearNot! is a virtual drama system that allows children to explore responses to and attitudes toward bullying behavior in the safety of a virtual environment (Aylett, Louchart, Dias, Paiva, & Vala, 2005). Children interact with the system as 'invisible friends' of a fictional character who is the victim of bullying and influence the character's behavior by suggesting possible actions the character should take in response. The system's fictional characters can sense and effect their environment, they have a model of emotion that influences their goal selection, and each character uses a partial order planner to find sequences of actions that accomplish internal goals. Characters also display unplanned reactive behavior when prompted by the environment. The reactive behavior allows characters to display emotional responses to unexpected acts of aggression that disrupt the user's suggested flow of events. The system's story arises from unscripted interactions between the characters and is guided by the user's prompts and suggestions.

Kelly, Botea, and Koenig (Kelly, Botea, & Koenig, 2007) show that planning can be integrated into large-scale commercial games to generate complex non-player character (NPC) behavior. The system works with the game *The Elder Scrolls IV: Oblivion*, where it generates NPC behavior using one or more HTNs. The system automatically converts plans it creates to solve game-specific planning problems into in-game scripts that can be used to control NPCs at run time. This system can be used to generate many NPC behavior patterns and automatically execute them in the game environment, potentially creating a higher level of immersion for players.

Finally, Cavazza, Pizzi, Charles, Vogt, and André (2009) developed a system that makes use of a narrative planner that reasons about emotional states in order to adapt three chapters of the novel *Madame Bovary* to drive an interactive experience. The user plays the role of the character Rodolphe and interacts with the virtual character Emma through natural language, spoken into a microphone. As the user speaks with the system, it maps the tone of the user's speech to a set of emotional categories. Once a user's speech act has been classified, the classification informs a Heuristic Search Planning

(Bonet & Geffner, 2001) (HSP) algorithm – a class of planning algorithms capable of returning plans quickly in a real-time system environment – that controls Emma's communicative actions. In this way, the narrative adapts quickly and responds emotionally to the speech acts chosen by the user.

4.2 Planning Plot

A second method of generating interactive narrative experiences is to create systems that ensure interaction with virtual characters result in a well-formed plot by coordinating behavior globally across all agents. *OPIATE* (Fairclough, 2004) operates according to this principle using plot operators taken from Propp's morphology (Propp, 1984), a taxonomy of schemata characterizing the structure of Russian folk tales. The system proceeds through a series of these operators and assigns goals to virtual characters by casting each into a story role based on their opinion of the user's character. The user affects the progression of plot functions and what role each character plays within the plot by influencing the opinion of the virtual characters.

LOGTELL (Karlsson, Ciarlini, Feijó, & Furtado, 2006) is a plot creation and visualization tool that combines planning, plan recognition, and user interaction in the story generation process. The system is capable of generating genre-specific plots using a planner and goal inference rules that govern the behavior of available character types. It also maintains a library of typical plans that can be used to quickly accomplish goals. A user interacts with the system through a plan visualization interface where they are able to alter the plan structure before sending the plot to the system's game engine for visualization.

Barber and Kudenko (2007) present a system that generates interactive plots that consist of dramatic dilemmas. Each dilemma serves as a goal for story plans produced by the system. The system allows its user to act as a character within a genre-specific story world, where the system presents a number of plot events before forcing the user to resolve a dilemma that their character is involved in. The system presents story actions as written text and allows the user to interact through simple text commands. The section of plot immediately following the resolution of a dilemma is influenced by the user's decision and the system plans the next section of plot while the human deliberates.

Sharma, Ontañón, Strong, Mehta, and Ram (2007) introduce a system that incorporates a model of player preference into the interactive plot generation process. The system uses a text-based interface for the interactive story game *Anchorhead* (Nelson & Mateas, 2005) which allows users to enter simple text commands to explore their environment. The system utilizes an expectation-maximization planning algorithm to select actions made by a drama manager (Weyhrauch & Bates, 1997), a system component that oversees the unfolding narrative and adjusts the structrure of the plot in response to user activity. These adjustments are intended to lead the story in an interesting direction based on the user's player profile.

As described in Section 2, Porteous, Cavazza, and Charles (2010b) use state constraints to specify a set of story trajectories that limit narrative plans produced by their system to those with interesting narrative properties. Porteous et al. (2010a) describe an interactive system that takes advantage of their constraintbased approach to narrative generation where the user can intervene during a narrative's visualization and change the state of the world to one unanticipated by the original system. They show that their current system runs fast enough to adapt and re-plan in real-time while conforming to the narrative constraints imposed on the system. Porteous et al. Porteous, Teutenberg, Charles, and Cavazza (2011) subsequently extended their work by including a more expressive temporal model in their knowledge representation.

Cavazza and Charles (2013) propose the study of interactive narrative medicine, a framework for dramatizing patient history for the benefit of physicians. Charles, Cavazza, Smith, Georg, and Porteous (2013) presents an interactive narrative system for patient education where a user experiences potential situations they will encounter during the clinical process in a virtual environment. Finally, Porteous, Charles, and Cavazza present NetworkING (Porteous, Charles, & Cavazza, 2013), a framework for generating plot structure by analyzing the shifting structure of social relationships within a group of people.

Mimesis (M. Riedl, Saretto, & Young, 2003) is an architecture for building and executing adaptive interactive narratives. The system consists of a 3D game engine, a *mediator* element that monitors user activity, and narrative planner. The system allows a user to act as a character within a story generated by the a narrative planner. However, the user is free to act out of character and may perform actions that change the world state in ways that the original plan did not anticipate. When this occurs, the mediator detects these deviations and is capable of responding to preserve story coherence, either by *intervening*, that is, effectively preventing the user's action from executing or by *accommodating*, that is, allowing the user to perform the intended action but then seamlessly creating a new plan where the user's action is consistent with the story line.

Extending the Mimesis mediation model, work by Harris and Young on *proactive mediation* (Harris & Young, 2009) integrates a plan recognition module into the system. This capability provides predictions about a user's likely plans he or she is pursuing and allows the system to make mediation decisions in anticipating of potential user deviations from a story. Instead of waiting for user actions to break the plot in order to intervene, the system can recognize the user's plan and modify the story and story world ahead of time in order to prevent the user's behavior from occurring. *Bidirectional accommodation* (Robertson & Young, 2013) extends the accommodation component of mediation by using a model of a player's knowledge of the story world and its history in order to consider a broader range of potential adjustments to story plans in the accomodation process.

The Automated Story Director (ASD) (M. O. Riedl, Stern, Dini, & Alderman, 2008) also extends the accommodation process of the Mimesis system. It introduces a four-tiered accommodation process that decreases search time during the re-planning process and introduces goal-substitution (a process by which alternative goals are swapped in to the story in place of ones that lead to conflicts in the accomodation process) when a solution to the initial problem cannot be found. Finally, the PAST system (Ramirez & Bulitko, 2012; Ramirez, Bulitko, & Spetch, 2013) combines ASD with PaSSAGE (Thue, Bulitko, Spetch, & Wasylishen, 2007), an interactive storyteller that learns a user's play style. This system is able to dynamically adapt a branching story to a user's play style as it learns.

5 Conclusion

A range of work has made use of plan-based knowledge representations and reasoning approaches to produce elements of narrative in both interactive and non-interactive environments. The close alignment between existing AI plan representations and both narrative theoretic and cognitive models of narrative structure has greatly facilitated work on the generation of story, discourse and interactivity. Narrative theoretic and cognitive/comprehension-focused models of narrative contain a range of elements that address issues beyond those paralleled by AI plan representations. For example, narratological concepts like unreliable narrators or focalization (Bal, 1997) are not specifically dealt with by planning models. Similarly, cognitive psychologists' models of salience in narrative understanding (Zwaan & Radvansky, 1998) are not found in plan generation approaches. Most AI approaches to narrative generation that leverage planning models use them as a base or a foundation, viewing them as the primitive buildings blocks that make up narrative structure. Additional work can then build on the primitive model to characterize more complex correlates. Examples of this type of approach include Bae and his collaborators' (2011) plan-based characterization of focalization and Cardone-Rivera and his colleagues' (2012) definition of Indexter, a plan-based model of salience in the context of story understanding.

Interestingly, the places where there are mismatches between conventional Ai representations of plans and narrative plans has also served as a point of success for plan-based narrative research, with efforts that have developed, for instance, specific models of conflict in story plans and explicit representations of and means for plan adaptation in the presence of user interactivity. This work has progressed successfully due in no small part to the well-founded semantics that AI knowledge representations provide and the underlying formal properties of the plan generation algorithms that create plans. While there are many structural and experiential features of narrative that are not yet addressed by planning-based models, future extensions, based similarly on well-founded semantics, hold the potential to increase our generation capability. Whether by proceeding in small, precise steps or by making large-scale representational changes, we're confident that plan-based methods for narrative generation hold a key to the capabilities for rich narrative content creation.

References

Abbott, H. P. (2008). The Cambridge introduction to narrative. Cambridge, England: Cambridge University Press.

Arijon, D. (1991). Grammar of the film language. Silman-James Press.

Arora, S., & Barak, B. (2009). Computational complexity: a modern approach. Cambridge, England: Cambridge University Press.

Aylett, R. S., Louchart, S., Dias, J., Paiva, A., & Vala, M. (2005). FearNot! – an experiment in emergent narrative. Intelligent Virtual Agents, 305–316.

Bae, B.-C., Cheong, Y.-G., & Young, R. M. (2011). Toward a computational model of focalization in narrative. In Proceedings of the 6th International Conference on the Foundations of Digital Games (FDG2011) (p. 313-315).

Bae, B.-C., & Young, R. M. (2008). A use of flashback and foreshadowing for surprise arousal in narrative using a plan-based approach. In Interactive Storytelling. Proceedings of the First Joint International Conference on Interactive Digital Storytelling, ICIDS 2008 (p. 156-67). Berlin, Germany.

Bal, M. (1997). Narratology: Introduction to the theory of narrative. University of Toronto Press.

Barber, H., & Kudenko, D. (2007). Dynamic generation of dilemma-based interactive narratives. In Artificial Intelligence and Interactive Digital Entertainment Conference (pp. 2–7).

Barber, H., & Kudenko, D. (2008). Generation of dilemma-based narratives: Method and turing test evaluation. In U. Spierling & N. Szilas (Eds.), Interactive Storytelling (Vol. 5334, p. 214-217). Berlin, Heidelberg: Springer.

Bonet, B., & Geffner, H. (2001). Planning as heuristic search. Artificial Intelligence, 129, 5–33.

Branigan, E. (2005). Projecting a camera: Language-games in film theory. CRC Press.

Brooks, C., & Warren, R. P. (1943). Understanding fiction. Prentice Hall. Bylander, T. (1991). Complexity results for planning. In J. Mylopoulos & R. Reiter (Eds.), Proceedings of the 12th International Joint Conference on Artificial Intelligence. Sydney, Australia, August 24–30, 1991 (pp. 274–279). Morgan Kaufmann.

Carbonell, J. G. (1981). Counterplanning: A strategy-based model of adversary planning in real-world situations. Artificial Intelligence, 16(3), 295–329.

Cardona-Rivera, R. E., Cassell, B. A., Ware, S. G., & Young, R. M. (2012). Indexter: A computational model of the event-indexing situation model for characterizing narratives. In M. A. Finlayson (Ed.), The Third Workshop on Computational Models of Narrative (pp. 32–41). İstanbul.

Cassell, B. A., & Young, R. M. (2013). Ember, toward salience-based cinematic generation. In Intelligent Narrative Technologies: Papers from the 2013 AIIDE Workshop. Palo Alto, CA, USA: AAAI Press.

Cavazza, M., & Charles, F. (2013). Towards interactive narrative medicine. Studies in Health Technology and Informatics, 184, 59–65.

Cavazza, M., Charles, F., & Mead, S. J. (2002a). Character-based interactive storytelling. IEEE Intelligent Systems special issue on AI in Interactive Entertainment, 17–24.

Cavazza, M., Charles, F., & Mead, S. J. (2002b). Planning characters' behaviour in interactive storytelling. The Journal of Visualization and Computer Animation, 13(2), 121–131. doi: 10.1002/vis.285

Cavazza, M., Pizzi, D., Charles, F., Vogt, T., & André, E. (2009). Emotional input for character-based interactive storytelling. In Autonomous Agents and Multi-Agent Systems (pp. 313–320).

Champandard, A., Verweij, T., & Straatman, R. (2009). The AI for Killzone 2's multiplayer bots. In Game Developers Conference (Vol. 2009, p. 4).

Charles, F., Cavazza, M., Smith, C., Georg, G., & Porteous, J. (2013). Instantiating interactive narratives from patient education documents. Artificial Intelligence in Medicine, 273–283.

Chatman, S. B. (1980). Story and discourse: Narrative structure in fiction and film. Cornell University Press.

Cheong, Y.-G., & Young, R. (2008). Narrative generation for suspense: modeling and evaluation. In Interactive Storytelling. Proceedings of the First Joint International Conference on Interactive Digital Storytelling, ICIDS 2008 (p. 144-55). Berlin, Germany.

Cheong, Y.-G., & Young, R. M. (to appear). Suspenser: A story generation system for suspense. IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games.

Christian, D. B., & Young, R. M. (2004). Comparing cognitive and computational models of narrative structure. In D. L. McGuinness & G. Ferguson (Eds.), Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative Applications of Artificial Intelligence, July 25-29, 2004, San Jose, California, USA (pp. 385–390). AAAI Press.

De Fina, A., & Georgakopoulou, A. (2012). Analyzing narrative: Discourse and sociolinguistic perspectives. Cambridge, England: Cambridge University Press.

Dehn, N. (1981). Story generation after TALE-SPIN. In IJCAI (Vol. 81, pp. 16–18).

Egri, L. (2009). The art of dramatic writing: Its basis in the creative interpretation of human motives. BN Publishing.

Elson, D. K., & McKeown, K. R. (2010). Tense and aspect assignment in narrative discourse. In Proceedings of the 6th International Natural Language Generation Conference (pp. 47–56).

Erol, K., Hendler, J. A., & Nau, D. S. (1994). UMCP: A sound and complete procedure for hierarchical task-network planning. In Proceedings of the Second International Conference on Artificial Intelligence Planning Systems (Vol. 94, pp. 249–254).

Fairclough, C. R. (2004). Story games and the OPIATE system (Unpublished doctoral dissertation). University of Dublin, Trinity College.

Fikes, R. E., & Nilsson, N. J. (1971). STRIPS: a new approach to the application of theorem proving to problem solving. Artificial Intelligence, 2, 189-208.

Gerrig, R. J. (1993). Experiencing narrative worlds: On the psychological activities of reading. Yale University Press.

Graesser, A., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. Psychological Review, 101(3), 371.

Gratch, J. (2000). Émile: Marshalling passions in training and education. In Proceedings of the 4th International Conference on Autonomous Agents (pp. 325–332).

Grice, H. P. (1989). Studies in the way of words. Cambridge, MA, USA: Harvard University Press.

Halliwell, S. (1987). Poetics of Aristotle: Translation and commentary. The University of North Carolina Press.

Harris, J., & Young, R. M. (2009). Proactive mediation in plan-based narrative environments. Computational Intelligence and AI in Games, 1(3), 233– 244.

Herman, D. (2007). Storytelling and the sciences of mind: Cognitive narratology, discursive psychology, and narratives in face-to-face interaction. Narrative, 15(3), 306–334.

Jenkins, H. (2004). Game design as narrative architecture. In N. Wardrip-Fruin & P. Harrigan (Eds.), First Person: New Media as Story, Performance and Game. Cambridge, MA, USA: MIT Press.

Jhala, A., & Young, R. M. (2010). Cinematic visual discourse: Representation, generation, and evaluation. IEEE Transactions on Computational Intelligence and AI in Games, 2(2), 69–81.

Karlsson, B., Ciarlini, A. E., Feijó, B., & Furtado, A. L. (2006). Applying a planrecognition/plan-generation paradigm to interactive storytelling. In ICAPS.

Kautz, H. (1998). The role of domain-specific knowledge in the planning as satisfiability framework. In Proceedings of the 4th International Conference on Artificial Intelligence Planning Systems (pp. 181–189).

Kelly, J.-P., Botea, A., & Koenig, S. (2007). Planning with hierarchical task networks in video games. In ICAPS-07 workshop on planning in games.

Lebowitz, M. (1985). Story-telling as planning and learning. Poetics, 14(6), 483–502.

Li, B., & Riedl, M. O. (2010). An offline planning approach to game plotline adaptation. In G. M. Youngblood & V. Bulitko (Eds.), Proceedings of the Sixth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE 2010, October 11-13, 2010, Stanford, California, USA. AAAI Press.

Louchart, S., & Aylett, R. (2003). Solving the narrative paradox in VEs – lessons from RPGs. Intelligent Virtual Agents, 244–248.

Loyall, A. B., & Bates, J. (1991). Hap: A reactive, adaptive architecture for agents (Tech. Rep. No. CMU-CS-91-147). Pittsburg, PA, USA: School of Computer Science, Carnegie Mellon University.

Magliano, J., Dijkstra, K., & Zwaan, R. A. (1996). Generating predictive inferences while viewing a movie. Discourse Processes, 22(3), 199–224.

Marsella, S. C., & Gratch, J. (2009). EMA: A process model of appraisal dynamics. Cognitive Systems Research, 10(1), 70–90.

Mateas, M., & Stern, A. (2003a). Façade: An experiment in building a fully- realized interactive drama (Vol. 2). Presented at the Game Developers Conference, Game Design track, San Jose, California.

Mateas, M., & Stern, A. (2003b). Integrating plot, character and natural language processing in the interactive drama façade. In S. Göbel, N. Braun, U. Spierling, J. Dechau, & H.

Diener (Eds.), Technologies for Interactive Digital Storytelling and Entertainment, TIDSE 03, proceedings (pp. 139– 151). Darmstadt: Fraunhofer IRB.

Meehan, J. R. (1977). TALE-SPIN, an interactive program that writes stories. In D. R. Reddy (Ed.), Proceedings of the 5th international joint conference on artificial intelligence. cambridge, ma, usa, august 1977 (pp. 91–98). William Kaufmann.

Moore, J. D., & Paris, C. L. (1993). Planning text for advisory dialogues: Capturing intentional and rhetorical information. Computational Linguistics, 19(4), 651-695.

Mueller, E. (2014). Computational modeling of narrative. Sprache und Datenverarbeitung, this volume.

Nelson, M. J., & Mateas, M. (2005). Search-based drama management in the interactive fiction anchorhead. In R. M. Young & J. E. Laird (Eds.), Proceedings of the First Artificial Intelligence and Interactive Digital Entertainment Conference, June 1-5, 2005, Marina del Rey, California, USA (pp. 99–104). AAAI Press.

Niehaus, J. M., Li, B., & Riedl, M. O. (2011). Automated scenario adaptation in support of intelligent tutoring systems. In FLAIRS conference.

Orkin, J. (2006). Three states and a plan: the AI of FEAR. Presented at the Game Developers Conference 2006.

Osborn, B. (2002). An agent-based architecture for generating interactive stories (Unpublished doctoral dissertation). Naval Post-graduate School, Monterey, CA, USA.

Paiva, A., Machado, I., & Prada, R. (2001). Heroes, villains, magicians,…: Dramatis personae in a virtual story creation environment. In Proceedings of the 6th international conference on intelligent user interfaces (pp. 129–136). New York, NY, USA: ACM.

Pizzi, D., Charles, F., Lugrin, J.-L., & Cavazza, M. (2007). Interactive storytelling with literary feelings. In A. C. Paiva, R. Prada, & R. W. Picard (Eds.), Affective Computing and Intelligent Interaction (Vol. 4738, p. 630-641). Berlin, Heidelberg: Springer.

Porteous, J., Cavazza, M., & Charles, F. (2010a). Applying planning to interactive storytelling: Narrative control using state constraints. ACM Transactions on Intelligent Systems and Technology, 1(2), 111–130.

Porteous, J., Cavazza, M., & Charles, F. (2010b). Narrative generation through characters' point of view. In Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (Vol. 1, pp. 1297– 1304). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Porteous, J., Charles, F., & Cavazza, M. (2013). NetworkING: using character relationships for interactive narrative generation. In M. L. Gini, O. Shehory, T. Ito, & C. M. Jonker (Eds.), International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '13, Saint Paul, MN, USA, May 6-10, 2013. IFAAMAS.

Porteous, J., Sebastia, L., & Hoffmann, J. (2001). On the extraction, ordering, and usage of landmarks in planning. In Proceedings of the Sixth European Conference on Planning (p. 174-182). AAAI Press.

Porteous, J., Teutenberg, J., Charles, F., & Cavazza, M. (2011). Controlling narrative time in interactive storytelling. In Autonomous Agents and Multiagent Systems (Vol. 2, pp. 449–456).

Propp, V. (1984). The theory and history of folklore. University of Minnesota Press.

Ramirez, A., & Bulitko, V. (2012). Telling interactive player-specific stories and planning for it: ASD + PaSSAGE = PAST. Presented at the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment. Ramirez, A., Bulitko, V., & Spetch, M. (2013). Evaluating planning-based experience managers for agency and fun in text-based interactive narrative. In Artificial Intelligence and Interactive Digital Entertainment Conference (pp. 65–71).

Riedl, M., Saretto, C. J., & Young, R. M. (2003). Managing interaction between users and agents in a multi-agent storytelling environment. In Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (pp. 741–748). New York, NY, USA: ACM.

Riedl, M. O., Stern, A., Dini, D. M., & Alderman, J. M. (2008). Dynamic experience management in virtual worlds for entertainment, education, and training. International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning, 23–42.

Riedl, M. O., & Young, R. M. (2005). An objective character believability evaluation procedure for multi-agent story generation systems. In

T. Panayiotopoulos, J. Gratch, R. Aylett, D. Ballin, P. Olivier, & T. Rist (Eds.), Intelligent virtual agents (Vol. 3661, p. 278-291). Berlin, Heidelberg: Springer.

Riedl, M. O., & Young, R. M. (2006). From linear story generation to branching story graphs. IEEE Computer Graphics and Applications, 23–31.

Riedl, M. O., & Young, R. M. (2010). Narrative planning: balancing plot and character. Journal of Artificial Intelligence Research, 39, 217–268.

Robertson, J., & Young, R. M. (2013). Modelling character knowledge in plan-based interactive narrative to extend accomodative mediation. In Intelligent Narrative Technologies (pp. 93–96).

Ryan, M.-L. (1991). Possible worlds, artificial intelligence, and narrative theory. Indiana University Press.

Ryan, M.-L., Herman, D., & Jahn, M. (2005). Routledge encyclopedia of narrative theory. Routledge.

Searle, J. R. (1969a). An essay in the philosophy of language. Cambridge, England: Cambridge University Press.

Searle, J. R. (1969b). Speech acts: An essay in the philosophy of language. Cambridge, England: Cambridge University Press.

Sgouros, N. M. (1999). Dynamic generation, management and resolution of interactive plots. Artificial Intelligence, 107(1), 29–62.

Sharma, M., Ontañón, S., Strong, C., Mehta, M., & Ram, A. (2007). Towards player preference modeling for drama management in interactive stories. In D. Wilson & G. Sutcliffe (Eds.), Proceedings of the Twentieth International Florida Artificial Intelligence Research Society Conference (pp. 571–576). AAAI Press.

Smith, T. C., & Witten, I. H. (1987). A planning mechanism for generating story text. Literary and Linguistic Computing, 2(2), 119–126.

Szilas, N. (2003). IDtension: a narrative engine for interactive drama. In S. Göbel, N. Braun, U. Spierling, J. Dechau, & H. Diener (Eds.), Technologies for Interactive Digital Storytelling and Entertainment, TIDSE 03, proceedings (Vol. 3, pp. 187–203). Fraunhofer IRB.

Thomas, J., & Young, R. M. (2010). Annie: Automated generation of adaptive learner guidance for fun serious games. IEEE Transactions on Learning Technologies, 3(4), 329-343.

Thue, D., Bulitko, V., Spetch, M., & Wasylishen, E. (2007). Interactive storytelling: A player modelling approach. In J. Schaeffer & M. Mateas (Eds.), Proceedings of the Third Artificial Intelligence and Interactive Digital Entertainment Conference, June 6-8, 2007, Stanford, California, USA (pp. 43– 48). The AAAI Press.

Tomaszewski, Z. (2011). On the use of reincorporation in interactive drama. In E. Tomai, D. Elson, & J. Rowe (Eds.), Intelligent Narrative Technologies IV: Papers from the 2011 AIIDE Workshop. AAAI Press.

Tomaszewski, Z., & Binsted, K. (2008). Marlinspike: An integrated approach to interactive drama. In H. Barber & D. Thue (Eds.), Integrating Technologies for Interactive Stories: Papers from the INTETAIN 2008 Workshop, January 7, 2008.

Trabasso, T., & Sperry, L. L. (1985). Causal relatedness and importance of story events. Journal of Memory and Language, 24(5), 595–611.

van den Broek, P., Pugles-Lorch, E., & Thurlow, R. (1996). Childrens' and adults' memory for television stories: The roles of causal factors, story-grammar categories, and hierarchical level. Child Development, 67, 3010-3028.

von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele. Mathematische Annalen, 100(1), 295–320.

Vorderer, P. (1996). Toward a psychological theory of suspense. In P. E. Vorderer, H. J. E. Wulff, & M. E. Friedrichsen (Eds.), Suspense: Conceptualizations, theoretical analyses, and empirical explorations. (p. 233-254). Lawrence Erlbaum Associates, Inc.

Walker, M. (1996). Limited attention and discourse structure. Computational Linguistics, 22(2), 255–264.

Ware, S. G., Young, R. M., Harrison, B., & Roberts, D. L. (2013). A computational model of narrative conflict at the fabula level. IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games. (forthcoming)

Ware, S. G., Young, R. M., Stith, C., & Wright, P. (2014). The best laid plans. In J. Orkin & N. R. Sturtevant (Eds.), Playable Experiences at AIIDE 2014 in proceedings of the 10th International Conference on Artificial Intelligence and Interactive Digital Entertainment (p. 205).

Weyhrauch, P., & Bates, J. (1997). Guiding interactive drama. Pittsburgh, PA: Carnegie Mellon University.

Young, R. M. (1999). Notes on the use of plan structures in the creation of interactive plot. In M. Mateas & P. Sengers (Eds.), Narrative Intelligence: Papers from the AAAI Fall Symposium (pp. 164–167).

Young, R. M., & Moore, J. D. (1994). DPOCL: A principled approach to discourse planning. In Proceedings of the Seventh International Generation Workshop (p. 13-20). Kennebunkport, ME.

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

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. Psychological Bulletin, 123(2), 162.