Computational Creativity in Narrative Generation: Utility and Novelty Based on Models of Story Comprehension

R. Michael Young

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

Introduction

Narrative is one of the fundamental modes of creative expression. Creativity plays a central role in the effective design of film, novels, oral stories, computer games and most other narrative media. Beyond the creation of narratives intended to entertain, narrative is used as a powerful mode of communication in a range of other contexts. Narrative and its creative design plays a role in education and training (Gee, 2003), the communication of scientific ideas (Hoffman, 2005), characterizations of business practices (Thomas, 1999), communicating cultural heritage (Collins et al, 2003) and other contexts. There is growing awareness of the foundational role of narrative in our understanding of not only fictional worlds (Gerrig, 1993) but also the real world around us (Bruner 1991).

While it is clear that the process of generating a narrative is systematic, and sometimes formulaic, the challenge of generating an interesting narrative is at the brink of our scientific and scholarly understanding (Graesser, Pomeroy, & Craig, 2001). That is, there is an insufficient explicit understanding of the means by which creativity shapes the design of a narrative artifact that has an interesting plot, an illuminating point, and a pacing that is captivating. We do know that much of the research on creativity shares a common notion of its nature: An artifact that is at once fulfilling a specific notion of *utility* while doing so in a way that is *novel* or unexpected (Perkins, 1981; Sternberg, 1997).

Computational approaches to creativity have often focused on methods for encoding the creation process as search through a design space (Lindsey, et al, 1980; Langley, et al, 1987). Rules or heuristics are used to encode notions of creativity by directing search through the design space in ways that focus the search on creative solutions or by modifying the search space to create designs that are unexpected or novel. Structural constrains on the space ensure that solutions found within it all meet the utility criteria, while directed search or careful re-definition of the space (Lenat 1983) give rise to solutions that satisfy the novelty constraint.

Recent approaches to the generation of narrative have employed similar search techniques. While these methods have shown initial successes, most are essentially focused inward, that is, they reason about a narrative's structure rather than how it will be experienced by its readers or viewers. Narrative is a communicative medium and much of its structure is composed by story tellers to change the mental state of those who experience it. A narrative is considered creative due directly to inferences made regarding its structure by viewers. A system seeking to build a creative narrative can exploit psychological models of narrative comprehension to produce stories that are both coherent (that is, they satisfy the utility constraint) and surprising (that is, they intentionally give rise to expectations that are later violated, thus satisfying the novelty constraint) (Graesser et al., 1998).

This paper describes very early work in progress whose objective is to develop new, cognitively informed computational models of creativity in the context of the generation of cinematic narrative – narrative that is told within 3 dimensional (3D) virtual environments. Motivated by narrative theoretic models of narrative structure (Polti, 1977; Chatman, 1980; Rimmon-Kenan, 1983) and psychological models of narrative comprehension (Graesser et al., 1998; Graesser & Klettke, 2001), we are developing techniques for creating stories and the cinematic techniques needed to convey them to users. These techniques will use these models to search for narratives that are at once coherent and surprising, satisfying these two constraints to generate creative narratives.

Related Work

The generation of creative cinematic narratives requires effective models of a broad class of narrative phenomena ranging from natural language generation (Cavazza and Charles, 2005) to character intention and affect (Mott et al, 2006). In the work we describe here, we focus on four central aspects of the process of narrative creativity: story generation, cinematic discourse generation, the modeling of user expectations and the combination of these three to create creative stories that are coherent and surprising.

Computational Models of Narrative Generation

Previous work on narrative generation systems can be broadly categorized into two types depending on the nature of the user's interaction within the story world environment. Non-interactive systems (Callaway and Lester 2002; Riedl and Young 2004)), do not allow the user to take action within the stories they create; rather, they focus on the automatic creation of narrative based on an (implicit or explicit) model of narrative production and their output has historically taken the form of narrative prose text. In contrast, the distinguishing feature of interactive narrative generation systems (Mateas and Stern 2005; Young and Riedl, 2003) is their focus on managing the adaptive flow of interaction for a user acting withing a story-based world.

Two projects that have focused on narrative generation have also taken creativity to be a core component. Bringsjord (2000) has created the BRUTUS story generator as a model of computational creativity. BRUTUS is constructed on the principle that wide variability in the possible output is akin to some form of creativity, and that this can be achieved through creating distinct and parameterizable components of the architecture for each substantial aspect of a story (such as plot, character, settings, or imagery). BRUTUS's processing first selects a thematic concept (a high level story schema) from a knowledge base. This "stage" is processed by a world simulator to instantiate each of the components and elaborate logical causes and effects for each event, creating a "plot". The "plot" is then parsed by a story grammar to determine which types of paragraphs and sentences will be used to tell the story, and finally, the sentences are created according to linguistic rules. BRUTUS's top-down, knowledge base driven generation places it in a large group of similar story generation systems, but it's emphasis on maximizing variability for the sake of creativity sets it apart.

Minstrel (Turner 1992) is a story generation system that uses case-based reasoning to implement a computational model of creativity. In Turner's system, story generation is modeled as a problem-solving task approached from the point of view of an author; problem solving in Minstrel uses four types of authorial goals: thematic goals, consistency goals, drama goals and presentation goals. The system begins with the top-level goal of telling a story, then decomposes its goals into sub-goals including adopting a theme, communicating the theme, applying drama goals, ensuring consistency, and achieving presentation goals.

Automatic Control of A Virtual Camera

For any 3D graphical system generating visual narratives, one of the fundamental operations is the manipulation of the system's viewport or camera. Many systems, like first person video games give user complete control over the positioning of the camera. In contrast, other systems (e.g., training simulations) greatly restrict the user's control over her viewpoint in order to ensure that the viewer gains specific knowledge of the world that satisfy the system's goals.

At the lowest level, manipulation of graphical viewport is achieved by selection of numerical values for the 6 degrees of freedom for the camera,namely x, y, z, pitch, roll, yaw, and other parameters of the camera like the FOV (Field of View). In order to create a cinematic experience, a system must determine the values for these variables for each frame of the cinematic artifact being produced. Several researchers have viewed this problem as one of constraint satisfaction over the geometric properties of the scene being filmed (Drucker and Zeltzer, 1995; Bares and Lester 1999).

One of the main limitations of geometric constraint solving systems used to generate cinematic narratives is their focus on the underlying scene geometry. Unfortunately, decisions they make about camera position do not do the distance in capturing the underlying reasoning about the cinematic motivation for many of the constraints. Film directors and cinematographers have developed effective visual storytelling techniques that can be cast as rules of thumb defining stereotypical ways of filming shot sequences in varying dramatic situations. These cinematic idioms have developed from practical experience; as audiences have learned to recognize these idioms, film-makers have come to rely on them as a means of communicating narrative structure. Several camera control architectures have been built to make use of cinematic idioms for virtual cinematography (Christianson et al. 1996).

Models of Narrative and Narrative Comprehension

Central to the work that we describe here is the modeling of the processes involved in narrative comprehension. In the planned research, we will adapt cognitive models of such processes to inform the narrative generation systems we build. By exploiting an explicit model of comprehension, our goal is to create interactive narratives whose structure and design are more readily understood by human users.

Research by cognitive psychologists strongly suggests that narrative is one of the fundamental ways in which we perceive, organize and understand our relationship to our environment (Bruner, 1986; Bruner, 1991; Graesser, Klettke, & Olde, 2001; Sarbin, 1986). Extensive research on the comprehension of narrative texts indicates that people's cognitive models of narrative are composed of plan-like networks of actions, temporally ordered and connected in a hierarchy by causal relationships between effects and preconditions. Directly analogous to these theoretical developments, our previous work (Christian and Young, 2004) has shown that there are strong correlations between cognitive structures used during narrative comprehension and the data structures built by AI planners (Penberthy and Weld, 1991; Young, Pollack and Moore, 1994).

Cognitive psychologists have identified six levels of structure used in the cognitive processes of narrative comprehension (Zwann and Radvansky, 1998). These are the surface code, explicit propositional textbase, situation model, thematic point, agent perspective, and genre. One of these levels, the situation model is particularly relevant to our work, containing elements most closely related to existing computational models of action and change (e.g., events in the plot, characters, settings, and the many complex relationships between these items). The situation model represents a viewer's mental model of the elements of a story and the relationships between them. The model is constructed incrementally during the reading of a narrative and can be broken down into a time-related series of individual models. Zwaan and Radvansky (1998) distinguish the (a) current model under construction, capturing information about the story conveyed in the most recently read phrase or sentence; (b) the integrated model comprising the series of models since the beginning of reading; and (c) the complete model, which is the last model obtained at the end of the narrative. When a new clause or sentence is read, a new current model is created to capture this new information and then incorporated into the integrated model in a process called updating.

Psychologists have identified five dimensions of situation models: space, time, causation, intentionality, and protagonists and objects (Graesser & Clark, 1985; Zwaan & Radvansky, 1998). The spatial dimension represents the physical layout of the imagined story world. The temporal dimension represents both the timeline in the story and the sequence in which the story is told. The dimension of causation maintains the causes and effects of events in the story, denoting which events cause others and why. The intentionality dimension tracks the goals of the characters in the story. Lastly, the protagonists and objects dimension follows the main character(s), lesser characters, and various items with which the characters come in contact (books, keys, etc.) (Zwann and Radvansky, 1998). These elements of the situation model all have correlates in the knowledge representations used by many planning systems. In the work we describe here, we are making these relationships explicit, using the generative capabilities of planners to create stories that make clear both the actions within the story and the effects on a user's situation model as she experiences the story's unfolding.

Researchers who investigate narrative comprehension have also attempted to extend their work from text to film (Graesser & Nakamura, 1982; Magliano, Miller, & Zwaan, 2001). For example, Graesser and Nakamura (1982) documented that people tend to notice and remember atypical actions in text that deviate from the normal scripts or schemas. For example, in a story about dining in a restaurant, they will remember whether or not the story mentions an unusual event (e.g., a customer falling on a floor) better than an irrelevant event (a customer reading a letter) and both of these better than a typical event (a customer eating). They cannot remember whether the typical events are explicitly mentioned or merely inferred. And what happens in films? Precisely the same pattern of data, except that the results are less pronounced. So people remember the atypical, irrelevant, and bizarre events more in text than in film, for reasons that are not perfectly well understood. More generally, it is currently not clear how well the results from text comprehension studies generalize to film. Graesser and Zwaan (1995) pointed out that the pragmatic ground rules are quite different in text, versus film, versus a videotape of everyday interactions. In film, we directly perceive spatial layout whereas this is often inferred in text; in text, we comprehend explicitly mentioned goals of characters, but these are normally inferred in film. Magliano et al. (2001) identified characteristics of the medium, such as sound effects, that diagnostically predict future events (as in suspense), whereas these features of media are absent in text.

Creative Cinematic Narrative Generation

Our current work addresses three key aspects of the generation of creative cinematic narrative. First, we are developing a model of inference and expectation during narrative comprehension. Second, we are building a model of cinematic discourse generation designed to direct that inference. Third, we are developing an algorithm that searches and constructs stories that violate expectation and give rise to creative experiences. This approach will explore the generation of story structure targeted at comprehension by adapting cognitive models of narrative comprehension for use within a plan-based story generation system operating within the context of a 3D camera control system capable of effectively conveying an underlying story using filmic idioms.

In addition to the formal models we are developing, each aspect of the research will be integrated with a system that explores the space of cinematic narratives and drives a commercial, off-the-shelf 3D game engine to render candidate cinematics. This work will build on the our experience constructing architectures for integrating research tools into 3D games (Riedl and Young, 2003; Young et al, 2004). While the knowledge gained by these engineering efforts will assist in the progress and evaluation of the work, the principal contributions of this work lies in the increased understanding of the computational models of narrative generation and the role these models play in the generation of creative narrative-based artifacts.

Story Generation Using Plan-Based Techniques

Narrative theorists broadly characterize narrative in terms of two constituent parts: the story and the discourse (Rimmon-Kenan, 2002). The story contains elements of the action within the story world, whereas the discourse contains elements of the telling of the narrative – the communicative actions which convey the story to the user. In our previous work, we have successfully employed

plan-based algorithms to generate low-level action sequences for story world plans. These story plans detail causal, temporal and intentional relationships between actions within a plot (Riedl and Young, 2005a; Riedl and Young, 2005b) and we have shown that they correspond to portions of cognitive models of story structure (Christian and Young, 2004).

Our planning approaches, however, have been limited by not demonstrating higher level narrative structure -- for instance, rising action, conflict of goals, reversal of fortune. We are combining our existing algorithms for generating fine-grained details for story world plans with narrative theoretic models of narrative structure (Polti 1977); our goal is to generate highly detailed and context sensitive plot lines that also are recognizably coherent in their global structure. We expect that the resulting narratives will be significantly more coherent and engaging than those we can currently produce.

To this end, we are exploiting the partial and hierarchical nature of the planning algorithm we employ. High-level descriptions of plot elements (such as Polti's (1977) set of dramatic situations) are represented as partially defined decomposition operators. Our planning algorithm is then used to build a hierarchically structured plot where constraints from these decomposition operators are propagated to lower level actions, constraining the global storyline appropriately. The resulting story planning algorithm will be used as a base for the creation of storylines used by the activities described in the following two sections.

Cinematic Control of Narrative

Within visual narrative media such as film, narrative discourse is typically composed of sequence of camera shots used to film the action unfolding within a story world. In this work, we focus on the role of a camera operating within a 3D virtual world and exploit parallels between the communicative nature of the camera and that of conventional textual discourse to develop generative models of cinematic discourse. Linguists commonly view natural language discourse as planned, communicative action (Grice, 1969). They model discourse structure based, in part, on the effects that each utterance is intended to have on the mental state of a reader (Moore and Paris, 1994; Grosz and Sidner, 1986). In this view, a single declarative sentence, for example, is intended to cause the reader to believe the truth of the proposition it asserts. Entire segments of text are intended to support the communicative goals of adjacent segments, for instance, when one paragraph provides background or justification for the claims of another (Mann and Thompson, 1987).

In this work, we are exploring the idea that 3D camera control in cinematic narrative can be viewed using the same intentional stance. In this model, cinematic discourse (i.e., multiple sequential camera shots filming elements of a single narrative) can be modeled in terms of the effects on the mental state of the viewer.

While there has been significant work on the use of planning techniques to determine the content and organization of textual discourse, there has been little work on the use of these techniques in the context of 3D camera control. In many ways, there is a natural fit when thinking of a single shot as a planned, intentional communicative act on the part of a cinematographer. Most shots have the effect of conveying to the viewer knowledge of the unfolding action captured within that shot's frame. Choices about the juxtaposition of two or more shots or shot sequences, similar to the construction of multi-sentential text, can signal intentions of the cinematographer relative to the meaning of the action portrayed in the shots.

Fortunately, the use of patterns of shots and shot sequences in film has been well-documented by cinematographers (Arijon, 1976). These film idioms are understood by cinematographers and film audiences as stereotypical ways of communicating specific story elements or contexts. In recent work (Jhala and Young, 2005a), we have begun the process of formalizing filmic communicative actions in terms of the actions' intentional constraints. This approach has been useful for selecting low-level camera operators for filming individual actions. In our current work, we are extending this representation to represent more abstract cinematic discourse structures intended to convey the internal structure of a cinematic discourse and integrate that with a cognitive model of inference described below.

Computational Model of Expectation and Surprise in Narrative

As discussed above, there have been a number of research efforts that have focused on the use of plan-based techniques for the creation of storylines for narrative systems. Most of this work has been concerned with the creation of stories whose properties meet criteria grounded within the story world (e.g., plan executability and soundness, character consistency, intentionality and coherence). Central to our work here is the idea that the generation of *creative* stories must also look outside the story, taking into account the effects the stories have on the mental state of the system's users as those stories play out.

In our research, we are integrating our existing plan-based techniques for creating and telling stories with cognitive models of narrative comprehension. The resulting algorithm explicitly models the inferences made by viewers of the stories that it creates. By making explicit reference to its models of viewer inference during the story construction process, the algorithm builds stories designed to promote targeted inferences about the stories' structure.

Specifically, we are developing models of the plan-based expectations that are formed by users as they experience our stories and use these models to build stories that exploit expectations to improve the user's understanding of key story elements. Specific aspects of comprehension that our project will target (e.g, memory, recall, study time on segments, engagement ratings) are

central to a user's interactive experiences in a range of contexts. For example, we might expect that there will be an increase in recall of causal relationships in the material when participants experience cinematics used in a narrative-based training simulation.

By constructing stories that take into account the updating process employed by users as they maintain their situation model, we seek to promote specific desired inferences about the structure of the story. In this manner, we parallel the approach taken by computational linguists working on the plan-based generation of natural language discourse (Moore and Paris 1994; Young, Moore and Pollack 1994). Although all five elements of the situation model have correlates within the knowledge representation schemes used by many planning algorithms, this work focuses on the development of a generative model that takes into account a subset consisting of causation, intentionality and protagonists and objects. While the remaining elements (i.e., space and time) clearly contribute to the process of narrative comprehension, the psychological evidence for modeling these factors is less amenable to computational formalization and we leave the pursuit of these aspects to follow-on research.

To model the updating process that users follow as they incrementally expand situation models, we will extend the knowledge representation used by our narrative planners to take into account the relations between story world and situation model. For example, each of the three elements of the situation model we are focusing on has explicit representational correlates within the narrative planning algorithms we currently employ. As is common in partial-order, causal link planners (Penberthy and Weld, 1991; Young, Pollack and Moore, 1994), causality is explicitly marked and reasoned about by causal links, annotations to our plans indicating when the effect of one action serves to establish a condition in the world needed as a precondition to a subsequent action's successful execution.

Intentionality is explicitly marked using frames of commitment (Riedl and Young, 2005) that provide intention-based justifications for actions performed by characters. Protagonists and objects are linked to schematized plan operators via each operator's variable list mapping instances of objects, characters, locations and other story world entities to the functional roles they play in the action's performance. While these representations have all been designed to fulfill other representational requirements, their semantics within the planning process mirror the semantics of the corresponding notions from the situation model. Schematized preconditions and effects that update the situation model will be incorporated into the operators used to build our stories. These will be used to track the dynamics of the situation model using the same plan construction processes we have already developed for narrative generation. In this manner, plans could be produced that have desired properties for the story world actions and the construction of the user's model of the story.

Due to the structure of narrative and the conventions of genre and culture, readers approach narratives with a collection of expectations about each aspect of the experience. Reader expectations drive, in part, the process of inference during comprehension; in turn, inferences made during comprehension shape a reader's subsequent expectations. While there are a number of distinct and sometimes mutually incompatible theories for inferences made during narrative comprehension (Allbritton, 2004; Rapp, McKoon and Ratcliff, 1992), we have adopted a restricted view of inferences – consistent with many psychologists' work – in which readers are modeled as problem solvers (Gerrig, 1993; Graesser et al., 1994) actively seeking to solve the plot-based problems faced by the characters within the story world. Specifically, we use a model of inference focusing on narrative comprehension as the construction of plans that underlie the story's progress. This computational approach to story comprehension suggests that expectations during comprehension arise from the need to integrate aspects of the user's situation model with the partially executed and partially articulated plans of the story's characters. Using the model of comprehension developed by the work described above, we are exploring the ways that stories can be automatically created and told in order to manipulate a user's plan-based expectations.

There are two major dimensions that predict recall of the actions, goals, plans, and events in narrative. One dimension is typicality. People tend to remember events and actions that are atypical rather than typical of a central organizing script, at least during short retention spans of an hour or so (Graesser et al., 1998; Graesser & Nakamura, 1992). The second dimension is coherence. People tend to remember information that is coherently related to other information in the text (Kintsch, 1998).

Given the capability to predict likely future-directed expectations from candidate narratives, our techniques will be used to explore the effective creation of expectations by the selection of stories prompting desired types of inferences about upcoming story action. Building on these techniques, we are investigating the creation of stories that foster specific expectations about upcoming story action and then violate those expectations in specific ways.

Planned Empirical Evaluation

The central goal of this research is to develop cognitively informed computational models of creativity in the context of narrative generation. The core notion in the design of these models is that the resulting cinematic narrative satisfies the utility/novelty criteria for creativity. Specifically, the model should produce cinematic narrative that is at once coherent while effectively manipulating users' expectations to create surprise. Our hypotheses focus on the notion that by designing narratives that are both coherent and surprising, viewers will perceive the stories as creative. To evaluate the work, we will test these hypotheses by 1) constructing an architecture that

automatically generates cinematics and 2) measuring the effect of our techniques on users' comprehension of the narratives they convey.

To conduct these experiments, we are first developing a narrative generation architecture using existing tools for extending existing commercial off-the-shelf game engines (e.g., Epic Software's Unreal Game Engine). In previous work, we have performed substantial research and experimental evaluation using these environments (Christian and Young, 2004; Riedl and Young, 2005; Young et al., 2004) and we will use tools developed as part of our previous efforts to facilitate the integration of this work into these narrative environments. Young has previously developed a set of web services that generate plans for activity within a virtual environment, then use those plans to drive characters and environmental changes within a computer game system. This architecture will be adapted to generate story structure and camera control directives for those stories and to then to record those actions as they play out within a game.

The resulting recordings will be saved as video files and presented to human subjects to evaluate various key aspects of their comprehension.

Summary

This project seeks to develop new, cognitively informed computational models of creativity in the context of the generation of cinematic narrative. The work has four major thrusts: 1) Developing a computational model of story generation that builds stories by integrating abstract and partial representations of story schemas with story plans produced by our previously developed story planning approaches, 2) Developing virtual cinematography models adapted from models of natural language discourse generation to control the 3D virtual camera in order to convey the actions of an unfolding narrative to a user, 3) Developing a computational model of inferences that arise during narrative comprehension, integrating cognitive models of comprehension with story and discourse creation, and 4) Developing a system that integrates these three parts to search for cinematic narratives that are both coherent and surprising.

References

Allbritton, D. W. (2004). Strategic production of predictive inferences during comprehension. Discourse Processes, 38, 309-322.

Arijon, D. (1976), Grammar of the film language. Silman-James Press..

Bares, W., and James C. Lester (1999). Intelligent multi-shot visualization interfaces for dynamic 3d worlds. In Proceedings of IUI, pages 119–126.

Bringsjord, S. and Ferrucci, D. (1999), *Artificial Intelligence and Literary Creativity: Inside the Mind of Brutus, a Storytelling Machine*, Hillsdale, NJ: Lawrence Erlbaum.

Bruner, J. (1991), The narrative construction of reality, *Critical Inquiry*, 18:1, 1-21.

Cavazza, M. and Fred Charles (2005), Dialogue Generation in Character-based Interactive Storytelling, In Proceedings of AIIDE 2005.

Callaway, C and Lester, J. (2002), Narrative Prose Generation. Artificial Intelligence, 139(2), pp. 213-252.

Chatman S. (1980), Story and discourse: Narrative Structure in Fiction and Film. Cornell University Press, Ithaca, New York.

Christian, D. B. and Young, R. M., (2004), Comparing Cognitive and Computational Models of Narrative Structure by. in AAAI 2004.

Christianson, D. B., Sean E. Anderson, Li-wei He, David H. Salesin, Daniel S. Weld, Michael F. Cohenn (1996), Declarative Camera Control for Automatic Cinematography, In AAAI 1996.

Collins, T.; Mulholland, P.; Bradbury, D.; Zdrahal, Z. (2003), Methodology and tools to support storytelling in cultural heritage forums, Proceedings. 14th International Workshop on Database and Expert Systems Applications, Issue , 1-5 Sept. 2003 Page(s): 105 – 109.

Drucker, S. M. and David Zeltzer (1995). Camdroid: A system for implementing intelligent camera control. In Symposium on Interactive 3D Graphics, pages 139–144.

Gee, J.P. (2003). What video games have to teach us about language and literacy. New York: Macmillan.

Gerrig, R. J. (1993). Experiencing narrative worlds. Yale University Press, New Haven, CT.

Graesser, A.C., & Clark, L.C. (1985). Structures and procedures of implicit knowledge. Norwood, NJ: Ablex.

Graesser, A.C., & Klettke, B. (2001). Plot, agency, and a structural affect theory of literary short story comprehension. In D. Schram and S. Steen (Eds.), Psychology and sociology of literature. (pp. 57-69). Amsterdam: John Benjamins.

Graesser, A. C., & Nakamura, G. V. (1982). The impact of schemas on comprehension and memory. In G. H. Bower (Ed.), The Psychology of Learning and Motivation, Vol. 16. New York: Academic Press.

Graesser, A.C., Pomeroy, V., & Craig, S. (2001). Psychological and computational research on theme comprehension. In M.M Louwerse and W. Van Peer (Eds.), Thematics in psychology and literary studies (pp. 19-34). Amsterdam: Benjamins.

Graesser, A. C., & Zwaan, R. A. (1995). Inference generation and the construction of situation models. In C. A. Weaver, S. Mannes, & C. R. Fletcher (Eds.), Discourse comprehension: Strategies and processing revisited (pp. 117-139). Hillsdale, NJ: Erlbaum.

Grice, H. P. (1969), Utterer's Meaning and Intention, The Philosophical Review 78: 147-77..

Grosz, B. and Candace L. Sidner. Attention, intentions, and the structure of discourse. *CL*, 12(3):175–204, 1986.

Hoffman, R. (2005), Storied Theory, Amercian Scientist, Vol. 93 Number 4 Page 308.

Jhala, Arnav and Young, R. Michael (2005), A Discourse Planning Approach for Cinematic Camera Control for Narratives in Virtual Environments, in AAAI 2005.

Kintsch, W. (1998). Comprehension: A paradigm for cognition. Cambridge, UK: Cambridge University Press.

Langley, P., Simon, H. A., Bradshaw, G. L. and Zytkow, J.M., Scientific discovery: computational exploration of the creative process. Cambridge, MA. MIT Press. 1987.

Lenat, D (1983). The role of heuristics in learning by discovery: Three case studies, In R.S. Michalski, J. G. Carbonel and T. M. Mitchell (Eds.), Machine Learning, an Artificial Intelligence Approach, (pp. 240 - 306), Palo Alto, CA. Tioga Press.

Lindsay, R., Buchanan, B., Feigenbaum, E., and Lederberg, J. (1980), DENDRAL, New York, McGraw-Hill.

Louwerse, M.M., & Van Peer, W. (2001)(Eds.), Thematics in psychology and literary studies. Amsterdam: Benjamins.

Magliano, J. P., Miller, J., & Zwaan, R. A. (2001). Indexing space and time in film understanding. Applied Cognitive Psychology, 15, 533-545.

Mann W. and Thompson S (1987). Rhetorical structure theory: A theory of text organization. USC/ISI Technical Report.

McKoon, G. and Ratcliff, R. (1992). Inference during reading. Psychological Review, 99, 440-466.

Moore, J. D. and Cecile Paris (1994). Planning text for advisory dialogues: Capturing intentional and rhetorical information. Computational Linguistics, 19(4):651–694.

Mott, B., Sunyoung Lee, and James Lester (2006), Probabilistic Goal Recognition in Interactive Narrative Environments, In AAAI 2006.

Penberthy, J. S. and Weld. D. (1991), UCPOP: A sound, complete partial order planner for ADL. In Proceedings of KRR, Cambridge, MA.

Perkins, D. (1981). The mind's best work. Cambridge: Harvard University Press.

Polti, G. (1977). Thirty-Six Dramatic Situations. The Writer Inc., Boston, USA., 1921; 1977.

Riedl, M. and R. Michael Young (2003). Character focused narrative planning. In Proceedings of the International Conference on Virtual Storytelling, Toulouse, France.

Riedl, M. and R. Michael Young. (2004a) A Planning Approach to Story Generation for History Education. Proceedings of NILE 2004.

Rimmon-Kenan, S. (1983). Narrative Fiction: Contemporary Poetics. Routledge, London.

Sarbin, T. R. (1986), Narrative Psychology. Praeger, New York.

Sternberg, R. J. (1997). Thinking styles. New York: Cambridge University Press.

Thomas, J. C. (1999) Narrative technology and the new millennium. Knowledge Management Journal, 2(9), 14-17.

Turner, S. (1992), MINSTREL, A computer model of creativity and storytelling, PhD dissertation, University of California, Los Angeles.

Young, R. M., and Moore, J. D. (1994a), DPOCL: A Principled Approach to Discourse Planning, In the *Proceedings of the Text Generation Workshop*, Kennebunkport.

Young, R. M., Johanna D. Moore, Martha E. Pollack (1994) Towards a Principled Representation of Discourse Plans, in the Proceedings of the Cognitive Science Society.

Young, R. Michael and Riedl, M. (2003), Towards an Architecture for Intelligent Control of Narrative in Interactive Virtual Worlds, In IUI 2003.

Young, R. M., Riedl, M., Branly, M., Martin, R. J., and Saretto, C. J. (2004), An architecture for integrating plan-based behavior generation with interactive game environments, *JOGD*, 1.

Young, R. M., Pollack, M.E., and Moore J.D. (1994), Decomposition and Causality in Partial-Order Planning. In *AIPS 2004*, Chicago, IL.

Zwaan, R.A., & Radvansky, G.A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, 123, 162-185.