

Comparing Effects of Different Cinematic Visualization Strategies on Viewer Comprehension

Arnav Jhala¹ and R. Michael Young²

¹ University of California, Santa Cruz, CA - 95064
jhala@cs.ucsc.edu

² North Carolina State University, Raleigh, NC - 27606
young@csc.ncsu.edu

Abstract. Computational storytelling systems have mainly focused on the construction and evaluation of textual discourse for communicating stories. Few intelligent camera systems have been built in 3D environments for effective visual communication of stories. The evaluation of effectiveness of these systems, if any, has focused mainly on the run-time performance of the camera placement algorithms. The purpose of this paper is to present a systematic cognitive-based evaluation methodology to compare effects of different cinematic visualization strategies on viewer comprehension of stories. In particular, an evaluation of automatically generated visualizations from Darshak, a cinematic planning system, against different hand-generated visualization strategies is presented. The methodology used in the empirical evaluation is based on QUEST, a cognitive framework for question-answering in the context of stories, that provides validated predictors for measuring story coherence in readers. Data collected from viewers, who watch the same story rendered with three different visualization strategies, is compared with QUEST's predictor metrics. Initial data analysis establishes significant effect on choice of visualization strategy on story comprehension. It further shows a significant effect of visualization strategy selected by Darshak on viewers' measured story coherence.

Keywords: Intelligent Camera Control, Computational Models of Narrative, Discourse Comprehension, Visual Discourse.

1 Introduction

The automatic generation and communication of narrative are long-standing research areas within Artificial Intelligence [15,11,13,4]. To date, much of the research on story generation has focused either on computational models of plot and the construction of story events or on the communication of a given story with text. Problems central to the work on textual communication of a narrative have much in common with challenges found in the generation of other genres of natural language discourse, including the critical issues of content determination (what propositions should appear in the text) and ordering (how should the propositions describing a story be ordered so as to present a coherent story to a reader).

Text-based storytelling systems (e.g., [4]) typically build upon computational models of discourse generation in order to produce coherent narratives. While these text-based

systems have been successful within the text medium, less attention has been focused on techniques for the creation of effective *cinematic* discourse – the creation of narrative told using a virtual camera operating within a 3D environment. As we know from our experiences as film-watchers, cinema is a powerful and effective medium for communicating narratives. In this paper, we describe an evaluation of computational techniques for constructing cinematic discourse, i.e. communication of stories through the visual medium. We evaluate the coherence of the story through indirect measurement obtained from viewers' judgements about the goodness of answers provided to question-answer pairs pertaining to stories viewed as cinematic sequences.

We present an evaluation of cinematic discourse generated by Darshak, an existing visual discourse planning system, with our novel experimental design. In Darshak [10,9] the cinematic conventions developed by film-makers are represented as action operators within a planning framework. These operators are used by a specialized planning algorithm that constructs visual discourse plans containing cinematic actions. Camera shots are represented as primitive operators that manipulate the beliefs of viewers about the state of the story world. Abstract operators are used to define patterns of storytelling and impose constraints on lower-level shots and shot sequences.

To evaluate the coherence of the cinematic discourse produced by Darshak, we present a detailed user evaluation that exploits techniques first used to evaluate QUEST [8], a cognitive model of question-answering in the context of stories. Results from our initial experiments show that choice of visualization strategy significantly affects viewer comprehension, and that cinematics generated by the system using an intentional model of communication result in improved comprehension over naive approaches. Our experimental design itself is a significant contribution toward cognition-based evaluation of cinematic narrative discourse generation systems.

A bipartite representation of narrative proposed by Chatman [14] is employed, describing narrative as containing both story and discourse. The story level includes the story world with all its content, characters, actions, events, and settings. The discourse level involves the telling of the narrative – the ordering of the events in the story chosen for recounting, and the linguistic communicative actions used to tell the story. In this paper, the focus is on the evaluation of automatically generated discourse level content, specifically narrative discourse that is communicated visually through the use of cinematic conventions.

One central problem in automatic generation of cinematic narrative discourse is the selection of viewpoints in 3D space that follows cinematic conventions and communicates the elements of the narrative unfolding in the 3D environment. Previous work on intelligent camera control has focused mainly on the graphical placement of the camera to satisfy given cinematic constraints [3,12,6]. Less attention has been paid on informing the placement of the camera based on the context of the story events [1,9]. Evaluation of these camera systems have either been based on runtime performance [7], or the efficacy of user-modeling based on specific interaction models in interactive scenarios [2]. Unlike other camera control systems, the Darshak system [9,10] incorporates an explicit representation of the story elements and constructs camera plans using plan operators that encode cinematic conventions. This paper presents an evaluation of the

output of the Darshak system with specific focus on measuring the coherence of the stories perceived by viewers on watching the movie clips automatically generated by the system.

2 Effective Cinematic Narrative Discourse

Within the context of our work, *cinematic narrative discourse* is a recounting of events occurring in a 3D graphical story world using a virtual camera. In this regard, events – and the world-state transitions that they prompt – are central to the notion of narrative. Linked to these events and state transitions are story world elements like settings, objects, characters and their internal beliefs, desires, plans, and goals. Film directors and master storytellers take advantage of these properties of stories and exploit them to craft interesting recountings or telling for these stories. Cinematics crafted by experts through manipulation of these properties are effective when viewers find the stories communicated by them coherent and the visualizations aesthetically pleasing. In order to produce a system that generates effective cinematics, we define three properties of the narrative, specifically saliency, coherence, and temporal consistency. We have designed a cinematic discourse generation algorithm to produce cinematics that demonstrate these properties in the narratives it produces. While these properties have not been explicitly addressed in previous approaches to the automatic generation of story visualizations, we claim that these properties are central to the comprehension of cinematic narrative.

Selection of salient elements: Salient elements in a cinematic discourse are elements from the story (e.g., events, characters, objects, the relationships between them) that are relevant to inferences needed for comprehension. When stories are narrated through any medium, it is the narrator’s responsibility to utilize the properties of the medium to maintain the engagement of the audience by providing them with the relevant information at the right times during the story’s telling. It is important for effective narrative discourse to maintain the focus of the audience on these salient story elements. Inclusion of extraneous elements that are not part of the causal chain of the story or leaving out necessary details of the story can interfere with the audience’s comprehension process and prevent them from enjoying the narrative experience. Choices made by a narrative generator regarding the content from the story to include in its telling directly affect salience and thus comprehension.

Plot coherence: Plot coherence can be described as the perception by the audience that the main events of a story are causally relevant to the outcome of the story. This definition of plot coherence is taken from the *Fabulist* story generation system [13].

In our work, plot coherence relates specifically to the the perceived understanding of the causal relationships between events that a viewer constructs during narrative comprehension. If events narrated to an audience seem unrelated to the story’s final outcome, then the audience may make incorrect inferences about the narrative and the narrative discourse may fail to achieve the author’s communicative goals.

Temporal consistency: In this work, temporal consistency refers to the consistency in the timing of the changes in a camera’s position and its movement in relation to the

events that are being filmed in the virtual world. Temporal consistency is important, as it is closely linked to the established conventions of cinematography readily identified by viewers. Any inconsistencies in the timing or placement of a camera affects the communicative meaning of the shot perceived by the audience. For example, moving the camera off its subject too early while filming an action leaves the viewer confused about the completion of the action and may introduce an unintended inference in the mind of the viewer. Viewers have grown familiar with many cinematic idioms that are routinely used to indicate directors' specific communicative intentions. It is thus important for a visual discourse generation system to be able to reason about appropriate temporal relationships between camera movement and story events executing in a 3D environment.

3 Visual Discourse Generation

Our approach for generating visual discourse is based on the Darshak system developed by Jhala and Young [9,10].¹ In this approach, cinematic discourse is generated by a hierarchical partial order causal link planner. The system takes as input an operator library, a representation of a story to be told, and a set of communicative goals relating to the story's telling. The operator library contains a collection of action operators that represent camera placement actions, transitions, abstract cinematic idioms and narrative patterns. Camera placement and transition actions, represented as primitive operators, have preconditions that encode continuity rules in cinematography and effects that alter a viewer's focus of attention. Operators representing abstract cinematic idioms and narrative patterns encode recipes for sequencing primitive or abstract operators and have effects that change the beliefs of a viewer about the story world and the actions in it.. The story to be told is input as a plan data structure that contains the description of the initial state of the story world, a set of story goals, and a totally ordered sequence of actions composing the story and the causal relationships between them. The input story plan is added to the knowledge base for the discourse planner in a declarative form using first-order predicates that describe the elements of the data structure, allowing the discourse-level plan operators to refer to story-level elements. The communicative goals are given to the system as a set of beliefs to be achieved in the mental state of the viewer.

The cinematic discourse planning algorithm performs both causal planning and temporal scheduling. To build a discourse plan, it selects camera operators from the operator library and adds them to the plan in order to satisfy specific communicative goals or preconditions of other communicative actions already in the plan. The algorithm binds variables in the camera operators, like a shot's start-time and end-time, linking them to corresponding actions in the story plan.

The output of the planning algorithm is a plan data structure containing a temporally ordered hierarchical structure of camera operators with all operator variables bound. This camera plan is combined with the story plan and is sent to an execution manager running on a game engine. The execution manager dispatches both story and camera

¹ Space limitations prevent a full description of the Darshak system. Readers are encouraged to see [9,10] for more detail.

actions on the game engine through objects that represent code for performing the action in the engine. In the game engine, story action function calls affect the physical state of the 3D world, such as movement of characters. Camera actions impose viewing constraints on the game world's virtual camera. A constraint-solving algorithm constantly checks the viewing constraints and maintains the camera's parameters (location and orientation) such that the constraints set by the camera actions are satisfied.

4 Empirical Evaluation

Evaluation of intelligent camera control systems is a challenging problem [7], primarily because there are several dimensions across which camera systems can be measured. Most current camera systems are evaluated based on their performance in terms of speed of calculating camera positions rather than their effectiveness at telling stories. While it is difficult to evaluate stylistic capabilities of such systems, it is possible to evaluate their efficacy in communicating the underlying narrative content. In this paper, we introduce an experimental design useful for comparing the effectiveness of different visualization strategies in communicating a story. Our approach is based on established cognitive models of story understanding [8] that have been successfully used to evaluate plan-based computational models of narrative [5,13].

To evaluate the efficacy of different visualization strategies, we prepared three visualizations of the same story, one with a fixed camera position within the setting, one with an over-the-shoulder camera following the protagonist, and one driven by a camera plan automatically generated by Darshak, a discourse planning algorithm [10]. Our purpose for running these experiments was two-fold: First, we want to investigate whether visualization strategies do indeed affect comprehension. Second, we sought to evaluate the quality of visualization generated by Darshak using a representation of camera shots as communicative actions. That is, we sought to determine whether visualizations generated by Darshak are coherent (as measured by viewers' perceptions of the attributes of the underlying stories). Empirical evaluation of such a subjective metric is challenging because a) viewers rarely share a common definition of coherence and so cannot be asked directly to give judgement on coherence, b) viewers differ in the relative magnitudes of values used to judge coherence, so the values they report cannot be directly mapped to a uniform scale across subjects, c) coherence is a property of both the cinematic discourse and the fabula plan itself, making the evaluation of the discourse difficult to separate from effects created by the underlying story. Any evaluation of the communicative elements must take into account this inherent coherence in the fabula itself, and d) it is difficult to control for the subjective opinions of subjects regarding coherence that are significantly influenced by extraneous factors such as the quality of character dialog or 3D animations.

In order to evaluate the effectiveness of various cinematic visualization techniques, we sought to measure how effective these techniques were at conveying the story structure that lies beneath the cinematic discourse. Because the underlying story elements in our system were defined as plan data structures themselves, we made use of previous work [5,13] relating these data structures to the mental models that users form during comprehension. To do this, we employ Christian and Young's mapping from plan data

structures onto a subset of the conceptual graph structures that model narrative defined by Graesser, et al, in their work on QUEST, a psychological model of question answering [8].

In the QUEST model [8] stories are represented as conceptual graph structures containing concept nodes and connective arcs. These graphs are called QUEST Knowledge Structures (or QKSs). They describe the reader's conception of narrative events and the relationships between them. Nodes and arcs in a QKS structure are based on their purpose in the narrative. For instance, if nodes A and B are two events in a story such that A causes or enables B, then A and B are represented by nodes in the QKS graph and are connected by a Consequence type of arc.

Techniques used by Graesser et. al. to validate the QUEST model were based on goodness-of-answer (GOA) ratings for question-answer pairs about the story shown to readers. In this approach, subjects read a short story, then were presented with a set of question-answer pairs relating to events in the story. For each pair, subjects were asked to provide a rating that measured how good the subject felt that the answer was an appropriate and accurate response to the question. GOA ratings obtained from their subjects were compared to ratings predicted by the QUEST model (QUEST supports questions of types why, how, when, enablement, and consequence). Graesser, et al's intent was to validate the QUEST algorithm as a model of human question-answering in the context of stories – the more closely QUEST's GOA ratings were to human subjects' responses over a large group of question-answer pairs, the more evidence was provided that QUEST's underlying question-answering model matched human performance. We make use of QUEST and its well-supported validation to gauge the mental models build by viewers as they watch cinematics, seeking to determine if these models capture specific relationships between events in a story.

Within the QUEST model, each event and goal is represented as a node in the QKS structure. The links in a QKS structure represent the different types of relationships between events and character goals within a story. Consequence(C), : The terminal event node is a consequence of the initiating event node. Reason(R), : The initiating goal node is the reason for the terminating event node. Initiate(I), : The initiating event node indicates a terminal goal node. Outcome(O), and : The terminal event node is the outcome of the initiating goal node. Implies(Im) are the types of relationship arcs between event and goal nodes in a QKS structure. : The initiating event node implies the terminal event node.

The algorithm for converting a POCL plan data structure to the corresponding QKS structure is shown in Figure 1. In our experiments, we first convert our story, represented as a plan data structure, into a corresponding QKS structure. Predictor variables proposed in the QUEST model can be used to calculate predictions for the goodness of answer (GOA) ratings - the measure of goodness of answer for a question/answer pair related to the events in the story. These GOA ratings are compared against data collected from participants who watch a video of the story filmed using static master-shots. The experiments measured the effectiveness of the QKS generated from a plan structure in predicting the Goodness of Answer (GOA) ratings given by the viewers.

GOA ratings in the models proposed by Graesser et.al. are determined through QUEST's predictor variables. The three predictors that are correlated to the GOA

Let n be the number of top-level goals of plan P and let m be the number of steps in P .

1. Create a total ordering o for the m steps in P that is consistent with the ordering constraints of P .
2. For each goal g_i of plan P for $i=1, \dots, n$, convert g_i into a goal node G_i .
3. For each plan step s_j for $j=1, \dots, m$ starting with the last step in o .
 - a. Convert s_j into a goal node G_j and an event node E_j
 - b. Link G_j to E_j with an outcome arc.
4. For each causal link in P connecting two steps s_1 and s_2 with condition c , connect the event node E_1 to the event node E_2 with a consequence arc.
5. For each causal link $\langle s_{1,p,q}, s_2 \rangle$ in P connecting two steps s_1 and s_2 , connect G_1 to G_2 with a reason arc.

Fig. 1. Algorithm for converting a POCL plan data structure to corresponding QKS

ratings are arc search, constraint satisfaction, and structural distance. Users who participated in our experiments were shown a video of a story through fixed viewpoints in a virtual world. They were then given question/answer pairs from the story and were asked to rate the quality of answers. These results were compared to the GOA ratings predicted by the QKS structure based on the underlying DPOCL plan. 15 participants were randomly assigned to three groups with different categories of questions of the forms: how, why and what enabled.

4.1 Method

Design. Two stories (S1 and S2) and three visualization strategies were used for each story (V1-fixed camera, V2-over-the-shoulder camera angle, and V3-Darshak driven camera) yielding 6 treatments. Here treatments are identified by labels with story label as prefix followed by the label of the visualization. For instance, S2V1 treatment refers to a visualization of the second story(S2) with fixed camera angle strategy (V1) Participants were randomly assigned to one of 6 groups (G1 to G6). Each participant was first shown a video and then asked to rate question-answer pairs of three forms of how, why and what enabled. The process was repeated for each subject with a second video.

For this experiment, 30 subjects were divided into Youden squares experimental design. Accordingly 6 subject groups of 5 subjects each were distributed across 6 treatments. This design was chosen in order both to account for the inherent coherence in the fabula and to account for the effects of watching several videos in order. Assuming a continuous response variable, the experimental design, known as a Youden square, combines Latin Squares with balanced, incomplete block design(BIBD). The Latin Square design is used to block on two sources of variation in complete blocks. Youden squares are used to block on two sources of variation - in this case, story and group - but cannot set up the complete blocks for latin squares designs. Each row (story) is a complete block for the visualisations, and the columns (groups) form a BIBDs. Since both group and visualisation appears only once for each story, tests involving the effects of visualisation are orthogonal for those testing the effects of the story type; The Youden square design isolates the effect of the visual perspective from the story effect. The stories for this experiment consisted of 15 steps corresponding to 70 QKS state/event-goal nodes. These numbers were chosen in order to keep the story lengths comparable to those used

Table 1. 2x3 Youden squares design for the experiment. G1 through G6 represent 6 groups of participants with 5 members in each group. They are arranged so that each story and visualization pair has a common group for other visualizations.

Viz	Master Shot	Over The Shoulder	Darshak
S1	G1,G4	G2,G5	G3,G6
S2	G5,G3	G6,G1	G4,G2

in earlier experiments [5,13]. The algorithm for converting plan data structures to QKS graphs is identical to that used in Christian and Young [5] and is shown for reference in Figure 1. The QKS graph of one of the stories used in the experiment is shown in Figure 2. Each story used in the experiment had 70 QKS nodes. Of the 70 QKS nodes, 10 and 12 questions were generated from randomly selected and converted to one of the three question types supported by QUEST: how, why, and what enabled. For each of the 10 questions, approximately 15 answer nodes were selected from nodes that were within a structural distance of 3 in the QKS graph generated from the story data structure. These numbers were chosen to have similar magnitude to the previous experiments, for better comparison. Each story used in the experiment had 70 QKS nodes. Of the 70 QKS nodes, 10 questions were generated from randomly selected QKS elements and converted to one of the three question types supported by QUEST: how, why, and what enabled. For each of the 10 questions, approximately 15 answer nodes were selected from nodes that were within a structural distance of 3 in the QKS graph generated from the story data structure. These numbers were chosen to have similar magnitude to Christian and Young’s previous experiments, for better comparison.

Procedure. Each participant went through three stages during the experiment. The entire experiment was carried out in a single session for each participant. Total time for a single participant was between 30 and 45 minutes. Initially, each participant was briefed on the experimental procedure and was asked to sign the consent form. They were then asked to read the instructions for participating in the study. After briefing, they watched a video of one story with a particular visualization according to the group assignment (Table 1). For each video, users provided GOA ratings for the question-answer pairs related to the story in the video. Participants were asked to rate the pairs along a four point scale (good, somewhat good, somewhat bad, bad). This procedure is consistent with earlier experiments [5,8] Next, they watched a second video with a different story and visualization followed by a questionnaire about the second story. The videos were shown in different orders to common groups in order to account for discrepancies arising from the order in which participants were shown the two videos.

5 Results

The mean overall GOA ratings recorded for the two stories are shown in Table 2 along with the standard deviations. These distributions of GOA scores do not present any problem for multiple regression analyses as the means do not show ceiling or floor




Fig. 2. KQS structure for one of the stories from the experiment. Goal and Event nodes from the story are represented by circles. Relationships between nodes is indicated by arrows which are labeled respectively as Reason(R), Consequence(C), and Outcome(O).


effects. The standard deviations are high enough to rule out the potential problem of there being a restricted range of ratings.

The GOA numbers shown in Table 2 indicate on preliminary observation that the GOA ratings for V1(Master Shot) and V3(Darhsak) are significantly closer than V2(Over-the-Shoulder shots). The standard deviations for V3 are lower than the other treatments in both stories. This indicates that participants converge better on rating questions in Darshak generated visualization. An interesting observation for V2 is that in story 2 the mean GOA ratings are significantly lower than the other two treatments with a significantly high standard deviation. These figures support the intuition that participants form their own interpretation of events in the story while looking at shots that are over-the-shoulder leading to the wide disparity in ratings in going from story 1 to story 2. While mean ratings provide an overall idea of the participant’s responses, it is interesting to observe disparity in GOA ratings for individual questions across different visualizations. Figure 3 summarizes mean GOA ratings for individual questions related to story 1 for the three visualization treatments. Question numbers 1, 8, and 10

Table 2. Mean GOA ratings and standard deviations from the experiment

GOA(stddev)	V1	V2	V3
S1	1.69 (0.91)	1.74 (0.82)	1.70 (0.79)
S2	1.76 (0.59)	1.51 (0.67)	1.78 (0.59)

are particularly interesting as there is a significant difference in the GOA ratings for the master shot visualization and the other two treatments, which have quite similar ratings. The question-answer pairs in discussion here are presented below: Why did Lane challenge Vinny? A. Because he wanted to kill Vinny. Why did Lane challenge Vinny? A. Because Lane wanted to steal tax money. Why did Lane meet Sheriff Bob? A. Because Lane needed a job. In Q1 and Q10 the ratings for V1 are significantly lower. This could be explained by examining the relationships between the question-answer nodes. In all three cases, the question answer nodes are two or more arcs away in distance along the causal chain of events. In case of the arc-search and structural distance predictors from QUEST these are good answers as they do lie on a causal chain of events leading to the question. The necessity and sufficiency constraints in the constraint satisfaction predictor reduce the strength of the answer. In Q1, for example, it is not necessary for Lane to challenge Vinny. He could just shoot him right away. This is an interesting case where users who were familiar with the gunfight setting chose to label the challenge as being an important step in killing Vinny. In a master-shot the gunfight sequence was not even recognized as a gunfight by most participants. Figure 4 shows the average ratings for each question for the second story. The interesting responses are the ones that have a significant difference in mean ratings across different visualizations. In this story, unlike story 1, the differences between ratings were relatively smaller. The interesting observations, however, were the ones where one of the treatments rated the answer as a 'bad' answer (rating < 1.5) and the other treatments rated the answer as a 'good' answer (rating > 1.5).

**Fig. 3.** GOA ratings for Story 1 across the three visualization strategies




Fig. 4. GOA ratings for Story 2 across the three visualization strategies

6 Conclusion and Future Work

Camera shots can be seen as intentional communicative actions used to manipulate the beliefs of viewers. In this paper we briefly described a representation of narrative patterns and cinematic idioms as plan operators and a reasoning algorithm for constructing visual narrative discourse. Further, we presented a novel evaluation technique to measure effectiveness of intelligent camera control systems to communicate a story and an experiment that used the methodology to evaluate our cinematic discourse generator. The evaluation is based on an established cognitive model of story comprehension. We also presented an experimental design to compare several different visualization strategies. Our initial results are encouraging. We found significant differences between GOA ratings obtained from participants viewing different visualizations of the same stories in support of our hypothesis that different visualization strategies did affect comprehension. We also found significant correlation between GOA ratings predicted by the QUEST predictors and two of the three visualization strategies.

These experiments serve towards achieving the long-term goal of defining cognitively plausible structures for representing story and discourse that a) can be generated automatically and b) is expressive enough to be used in procedural generation of narratives in rich virtual environments. Initial results described here indicate that a) different visualization strategies affect the viewer’s understanding of stories that have the same underlying fabula structure, and b) Visualization strategies selected by the discourse planning system, Darshak result in better comprehension as determined by the experiments used in evaluating the QUEST model. The work reported here is preliminary and has several limitations. While initial data analysis shows positive correlation with predicted GOA ratings, more data is needed to establish how visualization strategy affects comprehension for specific types of question-answer pairs (e.g. why, how, and what enabled). There are several limitations of work described in this paper. First, it is difficult to get subjects for experiments that contain several long videos and questionnaires. For this reason, the stories used for the experiments were short and the GOA

ratings were only solicited for a subset of question-answer pairs. Second, while there was consistency across the visualizations measure here, the effect of parameters like the visual appeal of the graphical engine, facial expressions and animations of characters, etc. needs to be studied further. This work does not include all the elements of the QUEST model as there is no direct mapping from plan structure to certain elements. Richer representation of story elements is needed for incorporating all the elements in the cognitive model. Future work will extend the plan representation to account for elements of the QUEST structure that do not feature in the current conversion algorithm. Further experiments are needed to isolate specific effects of presentation strategies on viewer comprehension.

References

1. Amerson, D., Kime, S., Young, R.M.: Real-time Cinematic Camera Control for Interactive Narratives. In: ACM ACE, Valencia, Spain, p. 369 (2005)
2. Bares, W., Lester, J.: Cinematographic User Models for Automated Realtime Camera Control in Dynamic 3D Environments. In: User Modeling, Vienna, Austria, pp. 215–226 (1997)
3. Bares, W., McDermott, S., Boudreaux, C., Thainimit, S.: Virtual 3D Camera Composition from Frame Constraints. In: ACM Multimedia, Los Angeles, USA, pp. 177–186 (2000)
4. Callaway, C.: Narrative Prose Generation. *Artificial Intelligence* 139, 213–252 (2002)
5. Christian, D., Young, R.M.: Comparing Cognitive and Computational Models of Narrative Structure. In: AAI, San Jose, USA, pp. 385–390 (2004)
6. Christianson, D., Anderson, S., He, L., Salesin, D., Weld, D., Cohen, M.: Declarative Camera Control for Automatic Cinematography. In: AAI/IAAI, Portland, USA, pp. 148–155 (1996)
7. Christie, M., Machap, R., Normand, J.-M., Olivier, P., Pickering, J.: Virtual Camera Planning: A Survey. In: Smart Graphics, Munich, Germany, pp. 40–52 (2005)
8. Graesser, A.C., Lang, K.L., Roberts, R.M.: Question Answering in the Context of Stories. *Journal of Experimental Psychology: General* 120(3), 254–277 (1991)
9. Jhala, A., Young, R.M.: A Discourse Planning Approach to Cinematic Camera Control for Narratives in Virtual Environments. In: AAI, Pittsburgh, USA, pp. 307–312 (2005)
10. Jhala, A., Young, R.M.: Representational Requirements for a Plan-Based Approach to Virtual Cinematography. In: AIIDE, Marina Del Rey, USA, pp. 36–41 (2006)
11. Meehan, J.R.: Tale-Spin, An Interactive Program that Writes Stories. In: IJCAI, Cambridge, USA, pp. 91–98 (1977)
12. Halper, N., Helbing, S., Strothotte, T.: A Camera Engine for Computer Games: Managing the Trade-off Between Constraint Satisfaction and Frame Coherence. In: Eurographics, Manchester, UK, vol. 20, pp. 174–183 (2001)
13. Riedl, M.O., Young, R.M.: An Intent-Driven Planner for Multi-Agent Story Generation. In: AAMAS, New York (2004)
14. Chatman, S.: *Story and Discourse: Narrative Structure in Fiction and Film*. Cornell University Press, Ithaca (1980)
15. Schank, R.C., Abelson, R.P.: *Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures*. L. Erlbaum, Hillsdale (1977)