# Planning Animation Cinematography and Shot Structure to Communicate Theme and Mood

Kevin Kennedy Cognitive Engineering Laboratory Department of Computer Science The University of Western Ontario London, Ontario, CANADA kevink@csd.uwo.ca

# ABSTRACT

Standard techniques, such as soundtrack recording, storyboarding and key-framing, are used to create animation adaptations of narratives. Many aspects of the narrative, such as moods, themes, character motivations and plot, must be captured in the audio-visual medium. Our work focusses on achieving the communication of moods and themes solely through the application of well-known cinematography techniques. We present a planning system that transforms a description of animator intentions and character actions into a series of camera shots which portray these intentions. The planner accomplishes this portrayal by utilizing lighting, framing, camera motion, colour choice and shot pacing. The final output is an animation that is intended to produce a viewer impression to support the animator's description of the mood and theme of the narrative.

# **Categories and Subject Descriptors**

I.2.1 [Computing Methodologies]: Artificial Intelligence— Applications and Expert Systems

## Keywords

Cinematography, computer animation, mood, theme, narrative, shot structure

# 1. INTRODUCTION

Creating an animation is a very difficult and time-consuming art. Traditionally, anything beyond a short animation has been beyond the capabilities of a single individual. The introduction of computers in this field has dramatically increased the capabilities of animators to create extended animations.

We have implemented a tool that helps animators to better capture a narrative in the visual medium, rather than to create entire animations without human intervention. The

Int. Symp. on Smart Graphics, June 11-13, 2002, Hawthorne, NY, USA. Copyright 2002 ACM 1-58113-555-6/02/0600...\$5.00

Robert E. Mercer Cognitive Engineering Laboratory Department of Computer Science The University of Western Ontario London, Ontario, CANADA mercer@csd.uwo.ca

computer is used as a knowledgeable assistant that applies rules of cinematography, such as lighting, colour choice, camera framing, montage, and camera angle, to the task of expressing the communicative goals of the human animator. This paper focusses on the shot structure and rendering functions within the context of our animation cinematography planning system.

Utilizing this system, an animator can take an existing animation and ask the computer to create a communicative plan that uses cinematography techniques to reinforce the content of the animation, or to create new meanings which would otherwise not be expressed. By harnessing cinematography information within a knowledge representation system, the computer can plan an animation as a presentation task and create cinematography effects on behalf of the animator. The system described in this paper contains a knowledge base of cinematography effects which is employed by a planner to automate the presentation of animations. The planner reasons directly about the communicative desizes that the animator wishes to express through the animation. The cinematography planning system can augment the visual vocabulary of an animator by acting as an expert assistant in the domain of cinematography technique. Our prototype implementation is not intended to be a complete animation creation environment. Rather it can be viewed as a potential addition to commercial animation production systems.

# 2. MOTIVATION

The computer is a tool with growing importance in the animation industry. Currently, computers are used to aid animators in performing some of the simpler repetitive tasks of generating images for animations. Research is steadily progressing in the direction of increasing the computer's role in the generation of animations.

Current (non-automated) animation techniques require either a large, skilled team of animators working together, or a single animator who devotes a great deal of time to generating a short animation. A sophisticated knowledge-based animation assistant could reason about an animation at the level of plot, theme, and character actions, motivations, and emotions, thereby enabling a single animator to create a compelling animation without concern for issues such as a detailed specification of low-level motion and geometry. Instead, the assistant would take responsibility for character visibility, action timing, and camera positioning, keeping in

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to distribute to lists, requires prior specific permission and/or fee.

mind the higher-level goals of the animator. Animation details could be easily changed in an interactive fashion without the animator expending much time or effort.

Such a complete system would be large and sophisticated. The system described in this paper implements a subset of the capabilities of a fully-featured animation assistant. It presents a novel interaction of computer animation with artificial intelligence (AI). To present their ideas and intentions, experienced animators can draw upon a large range of cinematographic techniques that are far more expressive than the simple presentation of spatial arrangements. These techniques include lighting, colour choice, framing, camera movement, and pacing. Our knowledge-based animation assistant aids animators in achieving their narrative goals by controlling the application of these presentation details.

# 3. PROBLEM OVERVIEW

## 3.1 Themes and Moods

Though the psychological effects of film techniques are hard to quantify, they are relentlessly sought by film-makers. By applying the proper visual effect, the film-maker can alter the emotional predisposition of a viewer. There is probably a certain level of learned response involved, as audiences are repeatedly exposed to the same stimuli in similar situations [22].

The broad approaches to setting mood are well understood. To set up a happy scene, one should use a technique such as bright lighting to create a "feeling up" mood. A sad, dramatic, or scary scene should use a low-key, "feeling down" effect such as cool colours. De-saturation can be used to draw viewers into the scene, or high saturation can make them feel like outside observers. These techniques are often applied in the visual medium to give films a specific mood.

Many themes also can be represented using cinematography. Often a director will use a certain colour symbol for the ongoing theme of a film. Other themes are more general and can be presented using a canonical film effect. For instance, a happy ending can be foreshadowed by using a lighting change from dark to bright. Good versus evil can be enhanced by associating good with brightness and colour, and evil with darkness and colourlessness.

Cinematography can also be used to give us insight into the mental state of a movie character. A common effect is to use a warped fish-eye camera lens to show mental instability. A startling zoom-in on a character helps us feel the terror that the character is also feeling. Sometimes a first-person-perspective moving camera is used to make us identify with the character who is fleeing for her life. All of these techniques are visual metaphors for psychological states that are easily understood by moviegoers.

## **3.2** Cinematography Techniques

To manipulate films in the way described above, the computer cinematographer must have a knowledge of cinematography and must be able to apply it. The types of cinematography knowledge our system contains is briefly described below. Further discussion of our representation and reasoning about this domain can be found in [17, 18].

## 3.2.1 Lighting

Lighting is used to set mood, direct viewer attention, and provide information. The computer cinematographer can apply lighting to characters and backgrounds independently. The quality of lighting can be adjusted to alter the amount and sharpness of shadows. The brightness and direction of lighting is changed to achieve communicative acts as required.

#### 3.2.2 Colour

The system described here has limited control of colour for scene objects. When object models are created they are also created with colour sets which can be applied to them before rendering. The colour sets fall into several overlapping classes of colour energy and colour saturation. The system can select a specific colour set which satisfies constraints imposed by the animator's communicative goals. The system does not contain a general model for the aesthetics of colour, but relies on the programmer to classify colours in terms of energy, temperature, and saturation.

#### 3.2.3 Camera Placement

The computer director of photography takes ownership of the virtual camera and its point of view. Given a scene containing background objects and characters, the system will orient the camera in a way that achieves the desired effects. The system presented here can only function with objects that are "well-behaved". The analogy used is that of a stage or small set. The computer can deal with objects that move through different positions on this small set, and arrive at proper camera placement solutions. An animation that involves large sweeping movements, interaction of convoluted objects, or highly constrained environments may not work correctly.

#### 3.2.4 Framing

Closely related to camera placement is the framing of objects within the two-dimensional field. When prompted by the director's communication goals, the computer will attempt to frame objects in certain zones of the screen to achieve corresponding visual effects.

#### 3.2.5 Shot Structure

In what is a step outside of the duties of a cinematographer, the system takes on some of the duties of a film editor. Given overall goals of pacing and rhythm, the computer will make decisions about where to place cuts in the film time-line.

To assemble an overall viewer impression of the scene environment, the computer will assemble short sequences of shots that portray important objects and relationships within a scene. The director can choose either an inductive or deductive approach to shot sequencing. The animator must also supply information about which objects and characters are important and what meaningful relationships exist between these objects.

# 4. RELATED WORK

## 4.1 Automating Visual Domain Presentations

This paper builds on techniques that have been developed in the domain of automated visual presentations. The research in this area has concentrated on presenting visual information to an interactive viewer using planning techniques. The communication being presented is usually for information, training, or advertising purposes. Seligmann and Feiner created a system to explain the operation of electronic equipment using annotated 3D illustrations [20]. The system took communicative acts as inputs and, as outputs, generated images of the operations being described. The system used design rules coupled with a generate and test approach.

André *et al.* described how their WIP system generated illustrated documents, using instructions for using a coffee maker as an example [1]. The system approached the problem as a layout task of presenting graphical explanations and textual descriptions together in the same 2D page. They reasoned about graphical constraints and layout issues. André and Rist described the use of Rhetorical Structure Theory (RST) as the planning paradigm that they used for generating their illustrated documents [2]. More recently, André and Rist have concentrated on using an animated presenter which interacts with the user [3, 4]. The presenter acts as an agent which helps to communicate the communicative tasks desired by the presentation's author.

Karp and Feiner tackled the problem of creating an animated presentation as a high-level planning task [15, 16]. They concentrated on a hierarchical decomposition of an animation presentation task into a sequence of cuts, dissolves and shots. The system used a heuristic, rule-based approach.

Bares and Lester created a system to automate animated explanations in a knowledge-based learning environment [7]. They are interesting in relation to this paper because they automatically place the camera to present information to the viewer, and to handle shot transitions. More recent work by this group has integrated natural language explanations with the animation in a coherent fashion [11].

## 4.2 Automating Cinematography Techniques

Early camera control research was concerned with controlling the camera based on what it was supposed to be viewing, as opposed to working with camera placement directly. Gleicher and Wilking built a system that allowed the user to control the camera by placing constraints on the image that the camera should produce [13]. Their technique allowed the user to interactively control the relations between world-space points and their image-space presentation. Their aim was to allow geometrically-complex camera motions to be specified easily and interactively.

Philips *et al.* created an automatic camera placement system that places and moves the camera to avoid obstructions, thus eliminating the need for complex camera movements [19]. The CINEMA system developed by Drucker *et al.* provides a language for users to specify camera movements procedurally [12]. Bares *et al.* [5, 8] use a real-time constraint solver to position a camera in a complex interactive 3D world so that the camera is positioned in a way that views characters and fulfills a cinematic goal. Another approach models the user's preferences to create a user-optimal camera placement to view a virtual 3D world [6]. In even more highly constrained situations, Bares *et al.* use a constraint solver to find a solution to various user-imposed camera viewing requirements [9], and Halper and Olivier use a genetic algorithm to find a good camera placement [14].

The two projects that could be considered most similar to this paper both involve the combined use of lighting and camera placement to create a visual presentation. Tomlinson *et al.* integrate a cinematography agent into a virtual animated agent environment [21]. This cinematography agent interacts with the agents in their environment and attempts to reflect the other agent's emotions using lighting changes and camera angles. Butz created a system that interacts with the PPP multimedia presentation planner (a successor to the WIP system discussed earlier) to create short animations to explain the function of technical devices [10]. The system uses visual effects, such as pointing the camera and shining spotlights, to convey a communicative goal.

### 4.3 Relation to Current Work

Like André et al., we have also found RST to be an appealing method for planning communicative acts. Our RST relations are used in the composition of communication structures that express theme and mood. Although similar, the works concerned with animation address different aspects of animated presentations. We view animations as a type of storytelling, and we provide a mechanism for directly reasoning about the communicative goals of the animator as artist. The approach of planning explicitly about communicative acts, and maintaining a knowledge base of cinematography techniques allows the system to reason in a uniform way about both the emotions of characters and the communication goals of the animator. Some of the research solves problems with camera placement that are not dealt with in this paper. Integrating these methods with the system described here would enhance its ability to handle more difficult camera placement requirements.

## 5. SYSTEM ARCHITECTURE

The system implemented in this paper has a somewhat traditional AI architecture consisting of a knowledge base, a planner and an acting agent, the latter being, in this case, a graphical renderer. The knowledge base contains knowledge about space, time, solid objects, lights, colours, cameras, scenes, shots, and cinematography effects. This knowledge is implemented using a language called LOOM. The planner creates a plan that implements the desired communicative acts expressed by the viewer. The planner itself is written in LISP. The renderer, also written in LISP, transforms the cinematography plan created by the planner into a sequence of animation frames which constitute the output of the program. The renderer makes use of the POVray ray-tracing system to create the graphical images.

Figure 1 shows the interactions between the various modules. The knowledge base acts as a source of information about cinematography, and as a store of knowledge while the animation plan is being generated. It is also used to store the planning rules that the planner assembles to create a solution. Hence the interaction between the planner and the knowledge base is bi-directional. The renderer, on the other hand, only retrieves information from the knowledge base; it does not add knowledge. There is also a direct connection from the planner to the renderer because the planner provides the renderer with the complete scene description that is to be rendered. User inputs are used to drive the various modules, and the renderer produces the final output.

#### 5.1 Tool or Assistant Approach

The computer cinematography assistant acts as a tool to be used by a human animator. The animator must do a great deal of the work involved with specifying the action, character blocking, and narrative. The human must also



Figure 1: Module Interactions

express her narrative goals in terms of the communicative acts that the system understands.

The computer will, however, assemble these elements, position the camera and lights, and generate the sequence of images that create the animation. Since this is meant as a semi-automated approach, the animator has the capability to overrule the computer and tell it to keep searching for a better cinematography "solution".

The human and computer work together as a team, just like the cooperation between a film director and her cinematographer on a movie set.

## 5.2 Input

In order to operate the program, a user must supply an animation description for the system to present. The animation description contains the following types of information:

- Characters. The names and initial locations of animation characters.
- Background Objects. The names and initial locations of background objects.
- Object Actions. Objects can be translated, stretched, squashed and rotated. Parabolic paths (vertical only) are also allowed for jumping. Any combination of these actions are allowed. Arbitrary curved paths (e.g. B-splines) are not supported. Characters do not have articulated limbs.
- Action Intervals. To help the program plan cuts, the beginning and end points of character actions should be supplied.

This animation description is a flat text file. The action description for the example animation discussed in Section 6 is as follows:

object "HouseSet", size: large, pos: stage-rear. object "Door", size: mid-size, pos: <0.6, 0, 2.5>.

character "SuperBall", size: small, pos: <0, 2, 0>. character "Pooch", size: small, pos: <-3, 0, -1>.

```
has-direction Pooch, rear-right.
action turn-left, time: 2 3, object: Door.
action move-forward, time: 3 5 char: Superball.
action turn-right, time: 5 6, char: Superball.
action move-rear-right, time: 6 9, char: Pooch.
action hop, time: 6 7, char: Pooch.
action hop, time: 7 8, char: Pooch.
action hop, time: 8 9, char: Pooch.
action tilt-back, time: 8 9, char: Superball.
action tilt-forward, time: 11 12, char: Superball.
```

After the user has created the animation description, she must tell the system which moods, themes, and information she would like expressed through the animation. These communicative acts are provided as a list to the planner (an example is given in Section 6). A graphical user interface is provided to facilitate the specification of these communicative acts.

#### 5.3 Knowledge Base

The knowledge base is a collection of information about the cinematography domain that is the basis for reasoning about animations. The knowledge base describes physical qualities, spatial relations, lights, colours, cameras, scenes, shots, and effects. This knowledge is used to assemble a plan about how to structure the cinematography of an animation. The knowledge base must be able to encode information about space, animations, and cinematography in a way that enhances reasoning about these domains.

The knowledge base is hierarchical in nature, that is, more complex concepts are built upon the definitions of simpler concepts. For example, a lightset is made up of individual spot-lights and fill-lights, which in turn are composed of various 1d-vectors and 3d-vectors which describe individual physical characteristics.

The knowledge base also makes use of qualitative representations whenever possible. Though many physical quantities are expressed with numerical valuations, they are always referred to at a higher level using their qualitative categorizations. For example, though small, mid-size, and large are defined as 0.5, 1.0, and 2.0 respectively, sizes are always referred to with small, mid-size, and large, never with their numerical valuations.

Further details can be found in previous papers [17, 18].

## 5.4 Shot Structure Creation

A "shot" is the sequence of images in an animation between any two camera cuts, that is, a single continuous motion captured by the virtual camera. In the knowledge base, the **shot** acts as a container for all of the elements needed to put a shot together. It defines camera targets, camera placement, and the beginning and end time of the film sequence. Camera shots are later spliced together to form a longer continuous film. Usually, the camera location is altered between camera shots to provide varying views of the action.

#### 5.4.1 Shotmaker Rule-Base

In order to fully exploit the cinematography techniques available to the RST planning system, a sequence of shots must be available to the planner to apply camera, framing, and lighting adjustments. The task of creating a shot sequence is carried out by a rule-based system called the shotmaker. The shotmaker analyzes the sequence of actions in the animation and produces an appropriate set of camera shots based on a small set of animator options.

The animator is able to control the pacing and the "sequence approach" of the shots produced by the shotmaker. The pacing controls the relative length of the shots. A fastpaced shot sequence has many short shots, while a slowpaced sequence depicts the same action with fewer, longer, shots. The sequence approach refers to the nature of the shots used to introduce the viewer to the subject matter of the animation. A deductive sequence approach begins with a long-shot comprising the entire setting which is followed by several close-up shots of individual characters. An inductive approach takes the opposite approach, beginning with character close-ups and progressing to more encompassing shots [22].

#### 5.4.2 Shotmaker Operation

The shotmaker approaches the animation as comprising three distinct phases: the introduction, the main action sequence, and the ending. The introduction takes place at the beginning of the animation and serves to introduce the viewer to the main characters and setting. The main-action phase consists of character actions that comprise the "plot" of the animation. Finally the ending consists of actions that take place at the end of the animation. Character actions that take place coincident with the beginning and ending of the animation are considered part of the introduction and ending phases respectively.

The shots that comprise the introduction are generated corresponding to the sequence approach chosen by the animator and the number of characters present in the animation. The characters and overall setting are shown in separate shots following the desired sequence approach. The shots for the main action sequence are decided by a different set of rules that consider shot pacing, and the temporal arrangement of actions. Several rules handle the cases in which actions are adjacent, overlapping, coincident or separate. Depending on the action pacing, action combinations can be arranged into different shot sequences. For example, in a fast-paced presentation, two overlapping actions are shown with 3 different shots: one showing the start of the first action, a second showing the overlap of the two actions, and a final shot showing the end of the trailing action. A slow-paced presentation would simply show both overlapping actions with one continuous shot.

The rule base only considers three aspects of the animation when creating the shot structure: action ordering, pacing, and sequence structure. The shotmaker is intended to be extensible to handle more influences on shot structure, such as character intentions and relationships. The output of the shotmaker is an ordered sequence of shots complete with timing information and camera-target choices. This information is used by the RST planner to create appropriate cinematography effects for each shot.

# 5.5 RST Planner

Rhetorical Structure Theory (RST) is the basis for reasoning about communicative acts. The planner is given an animation description and a set of communicative acts to achieve. The planner must generate a plan that presents these communicative acts in a direct and coherent way. Because the planner uses RST structures and because the planner can only generate coherent rhetorical structures, the generated plans are coherent.

The planner constructs RST plans which consist of cinematographic instructions for presenting animation scenes. The planner is a depth-first forward chainer that actively analyzes the effects of the RST plan steps. While the RST plan is being constructed, the planner searches through the space of all possible RST plans implied by the predefined RST plan steps. The partial RST plan at any point is the "state" of the planner as it searches through possible plans. Operation of the planner is described in greater detail elsewhere [18].

The planner transforms user-specified communicative acts into knowledge base assertions about the current **shot**. As described earlier, the **shot** is a knowledge base concept that contains information about how to render a single camera shot.

#### 5.6 Renderer

The rendering subsystem digests a **shot** instance and produces an input file for the ray-tracer. The renderer is a LISP program which is passed the name of the **shot** to be rendered. The renderer then queries the knowledge base to determine the details of what must be rendered in the **shot**. All necessary information about how to create the final visual images are contained in the **shot** instance.

Specifically, the renderer queries the knowledge base to find out which spatial objects are present, where the camera is located, and what light sources are illuminating the scene. The renderer must also animate the characters' movements according to the input animation's specifications.

The rendering subsystem must perform many calculations to locate the lights, camera, and objects correctly in relation to each other. All spotlights must be correctly pointed towards their targets, all fill lights must be oriented so that their faces lie perpendicular to their targets as well. The absolute camera position must be derived from the relative positions contained within the **shot** instance.

Another task faced by the renderer is how to correctly time the duration of the shot. The renderer can start and end a shot at any "real" time (real here meaning the time from the user-specified animation). This allows the renderer to position a shot anytime within the time sequence of the initial animation. By changing the number of frames rendered within the shot, the actual run time of the shot can be altered to create a fast-motion or slow-motion effect.

The ray-tracer used for this project is the Persistence of Vision ray-tracer, also known as POVray. A ray-tracer is a program that will produce a graphic image of a scene when given the set of objects, lights, and atmospheric effects. Raytracers are notoriously slow for rendering, and never used for real-time rendering, but they produce very detailed and accurate lighting, shadows, and reflections. Ray-tracing is required by our system because of the reliance on subtle lighting to produce visual effects.

The rendering sub-system encodes some aspects of cinematography knowledge. Knowledge that is not needed for planning communicative acts is implemented at this low level. This knowledge is procedural knowledge that is encoded directly by the LISP rendering program. Some examples of knowledge encoded by the renderer are:

• always film relative to the eye level of characters.

- do not "cross the line" and shoot from behind the set.
- expand light-set size when showing a long-shot.

The output of the rendering subsystem is a text file completely describing the shot for the ray-tracer. The renderer combines many object-model files, together with lighting and camera positioning information, into a final file which is passed to the ray-tracer. The final images of the ray-trace are shown to the animator for their approval.

# 6. AN EXAMPLE ANIMATION

The operation of the entire system can be shown with an example animation. Consider the following scenario as an animation. The hero, Superball, is standing outside a door. He opens the door and enters and is immediately set upon by a dog ("Pooch") who jumps toward him. There are two possible stories that could be interpreted from these events. The first is that Superball is arriving home and is being greeted by the family dog. The second is that Superball is entering a strange house and being attacked. The following example shows how this sequence of events can be presented in two ways each creating a separate version of the events. The second more "scary" version is of primary interest here because it makes more use of cinematography as communication.

Initially the system is provided with a physical description of the setting, characters, and actions. This is passed to the shotmaker to generate a suitable sequence of shots for the animation. In the case of the attack scenario, fast pacing is requested which results in a sequence of 8 shots.

Next the planner is invoked and is given three communicative acts that produce a scary scenario:

- Elaborate scene is scary
- Restatement Pooch has hidden identity
- Restatement Superball is frightened

The planner considers these goals and generates an RSTbased plan that achieves these communicative acts by using five separate cinematographic effects:

- Use close-up camera placements
- Use harsh lighting (with dark shadows)
- Use high-saturation, low-energy colours
- Place lights behind Pooch
- Camera-zoom towards Superball

The RST chart for this plan is shown in Figure 2. This RST plan is generated automatically by the system. The plan for the "happy" scene is much simpler, relying on bright lights and warm saturated colours.

Figures 3 through 7 show several frames from the two different animations. Frames from the "happy" animation are on the left and the "scary" animation frames are to the right. The frames show how the cinematography effects are applied to the same animation description to produce two different narrative interpretations of the same events. Happy shots on the left use medium-distance camera positions with high-key lighting and less shading. The scary frames show how darker lighting and close-up camera positioning create a more moody feeling. Figure 3 shows how rear lighting is used to conceal Pooch's identity, while Figure 5 shows a frame from a dramatic zoom-in to Superball's face to reflect fear.

Taken together, these two sets of frame excerpts show how the system can alter the animations mood and even meaning through correctly applied cinematography effects.

## 7. CONCLUSIONS AND FUTURE WORK

We have presented a knowledge-based system that automates the communication of theme and mood by an animator adapting a narrative to computer animation. In this presentation we have given an overview of an animation assistant focussing on the shot structure and rendering functions. We see this project contributing to a larger movement towards more intelligent computer animation tools. Implementing a practical animation production environment will require the involvement of artists skilled in the creation of animations.

Immediate enhancements that we are considering include:

- Parameterization of object actions so that the manipulation of character "acting" can be achieved in a limited but useful way.
- Integration with a more robust geometric camera positioning system would enhance the usability and generality of the system.
- More advanced shot structure arrangements need to be investigated to improve the expressive power of the animation system.

### 8. **REFERENCES**

- E. André, W. Finkler, W. Graf, T. Rist, A. Schauder, and W. Wahlster. WIP: The Automatic Synthesis of Multimodal Presentations. In M. T. Maybury, editor, *Intelligent Multimedia Interfaces*, pages 75–93.
   American Association for Artificial Intelligence, 1993.
- [2] E. André and T. Rist. The Design of Illustrated Documents as a Planning Task. In M. T. Maybury, editor, *Intelligent Multimedia Interfaces*, pages 94–116. American Association for Artificial Intelligence, 1993.
- [3] E. André and T. Rist. Coping with Temporal Constraints in Multimedia Presentation Planning. In Proceedings of the Thirteenth National Conference on Artificial Intelligence and the Eighth Innovative Applications of Artificial Intelligence Conference, volume 1, pages 142–147. AAAI Press/The MIT Press, 1996.
- [4] E. André and T. Rist. Presenting Through Performing: On the Use of Multiple Lifelike Characters in Knowledge-Based Presentation Systems. *Knowledge-Based Systems*, pages 3–13, March 2001.
- [5] W. H. Bares, J. P. Gregoire, and J. C. Lester. Realtime Constraint-Based Cinematography for Complex Interactive 3D Worlds. In *Proceedings of the Tenth National Conference on Innovative Applications* of Artificial Intelligence, pages 1101–1106. American Association of Artificial Intelligence, July 1998.
- [6] W. H. Bares and J. C. Lester. Cinematographic User Models for Automated Realtime Camera Control in



Figure 2: RST Plan for "Scary" animation





Figure 3: Happy Pooch and Scary Pooch





Figure 4: Superball Makes His Entrance





Figure 5: Camera Zoom Heightens Drama in Scary Animation











#### Figure 7: End of Scene

Dynamic 3D Environments. In Proceedings of the Sixth International Conference on User Modelling, pages 215–226, June 1997.

- [7] W. H. Bares and J. C. Lester. Realtime Generation of Customized 3D Animated Explanations for Knowledge-Based Learning Environments. In Proceedings of the Fourteenth National Conference on Artificial Intelligence, pages 347–354. American Association of Artificial Intelligence, July 1997.
- [8] W. H. Bares and J. C. Lester. Intelligent Multi-Shot Visualization Interfaces for Dynamic 3D Worlds. In Proceedings of the 1999 International Conference on Intelligent User Interface, pages 119–126, 1999.
- [9] W. H. Bares, S. Thainimit, and S. McDermott. A Model for Constraint-Based Camera Planning. In Smart Graphics: Papers from the 2000 AAAI Symposium, pages 84–91. AAAI Press, 2000.
- [10] A. Butz. Anymation with CATHI. In Proceedings of the 14th Annual National Conference on Artificial Intelligence (AAAI/IAAI), pages 957–962. AAAI Press, 1997.
- [11] B. H. Daniel, W. H. Bares, C. B. Callaway, and J. C. Lester. Student-Sensitive Multimodal Explanation Generation for 3D Learning Environments. In Proceedings of the Sixteenth National Conference on Artificial Intelligence, pages 114–120. American Association of Artificial Intelligence, July 1999.
- [12] S. M. Drucker, T. A. Galyean, and D. Zeltzer. CINEMA: A System for Proceedural Camera Movements. In *Proceeedings 1992 Symposium on Interactive 3D Graphics*. ACM SIGGRAPH, March–April 1992.
- [13] M. Gleicher and A. Witkin. Through-the-Lens Camera Control. *Computer Graphics*, 26(2):331–340, July 1992. SIGGRAPH '92.
- [14] N. Halper and P. Olivier. CAMPLAN: A Camera Planning Agent. In Smart Graphics: Papers from the

2000 AAAI Symposium, pages 92–100. AAAI Press, 2000.

- [15] P. Karp and S. Feiner. Issues in the Automated Generation of Animated Presentations. In *Proceedings Graphics Interface '90*, pages 39–48. Canadian Information Processing Society, May 1990.
- [16] P. Karp and S. Feiner. Automated Presentation Planning of Animation Using Task Decomposition with Heuristic Reasoning. In *Proceedings Graphics Interface '93*, pages 118–127. Canadian Information Processing Society, May 1993.
- [17] K. Kennedy and R. E. Mercer. Using Cinematography Knowledge to Communicate Animator Intentions. In Proceedings of the First International Symposium on Smart Graphics, pages 47–52. ACM Press, 2001.
- [18] K. Kennedy and R. E. Mercer. Using Communicative Acts to Plan the Cinematographic Structure of Animations. In *Proceedings of the Fifteenth Canadian Conference on Artificial Intelligence*, page to appear. Springer-Verlag, 2002.
- [19] G. B. Phillips, N. I. Badler, and J. Granieri. Automatic Viewing Control for 3D Direct Manipulation. In *Proceeedings 1992 Symposium on Interactive 3D Graphics*. ACM SIGGRAPH, March–April 1992.
- [20] D. D. Seligmann and S. Feiner. Automated Generation of Intent-Based 3D Illustrations. volume 25, pages 123–132. ACM SIGGRAPH, July 1991. SIGGRAPH '91 Conference Proceedings.
- [21] B. Tomlinson, B. Blumberg, and D. Nain. Expressive Autonomous Cinematography for Interactive Virtual Environments. In *Proceedings of the Fourth International Conference on Autonomous Agents*, pages 317–324. ACM Press, June 2000.
- [22] H. Zettl. Sight Sound Motion: Applied Media Aesthetics. Wadsworth Publishing Company, 1990.