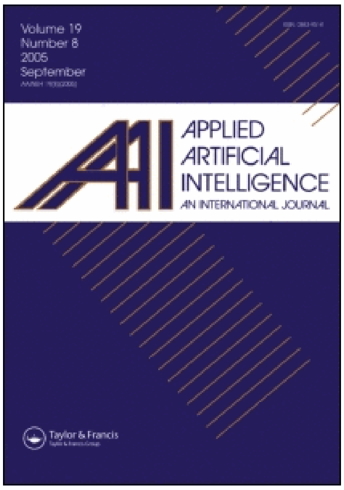

Applied Artificial Intelligence
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713191765

EMOTIONS FOR A MOTION: RAPID DEVELOPMENT OF BELIEVABLE PATHEMATIC AGENTS IN INTELLIGENT VIRTUAL ENVIRONMENTS

Carlos A. Martinho; Ana M. Paiva; Mario R. Gomes

To cite this Article Martinho, Carlos A. , Paiva, Ana M. and Gomes, Mario R.(2000) 'EMOTIONS FOR A MOTION: RAPID DEVELOPMENT OF BELIEVABLE PATHEMATIC AGENTS IN INTELLIGENT VIRTUAL ENVIRONMENTS', Applied Artificial Intelligence, 14: 1, 33 — 68

To link to this Article: DOI: 10.1080/088395100117151

URL: http://dx.doi.org/10.1080/088395100117151

Applied Artificial Intelligence, 14:33–68, 2000

0883-9514/00 $12.00 + .00

EMOTIONS FOR A MOTION: RAPID DEVELOPMENT OF BELIEVABLE PATHEMATIC AGENTS IN INTELLIGENT VIRTUAL ENVIRONMENTS

CARLOS A. MARTINHO, ANA M. PAIVA, and MÁRIO R. GOMES
Instituto de Engenharia de Sistemas e Computadores, Lisboa, Portugal

This paper discusses and proposes a solution for two interrelated problems: how to develop real-time intelligent virtual environments with believable agents over short periods of time, and how to instill emotional believability in these agents. Short development time is achieved using a modular architecture approach, minimizing the communication between modules, and therefore allowing their parallel implementation. Real-time performance is achieved using information redundancy and one-way communication between modules. All these techniques are supported by a theatrical metaphor. The intelligent virtual environment is divided into three functional units: the dynamic scriptwriter who creates the narrative, the theatrical company and its cast of actors who perform upon it, and the virtual stage managers who handle the cameras and special effects. The dynamic scriptwriter is responsible for the management of the synthetic actor inner persona. Personality is expressed as a regular exteriorization of emotions. Thus, the synthetic characters are characterized by a set of valenced reactions exteriorizing their inner goals, beliefs, and attitudes. The emotional structure is handled by an instantiation of Ortony, Clore, and Collins' cognitive theory of emotions and the agent implementation is inspired by Frijda's two-stage emotional stimuli appraisal theory.

This approach was tested in a real-time application developed under a tight schedule, and exhibited in Lisbon during the last world exposition of the century, Expo'98. Two synthetic dolphins, the emotional characters of this intelligent virtual environment, were displayed to more than one million visitors over the four months of the exhibition. This

This work was supported by PRAXIS XXI BM/12597/97.

We would like to express a special thank you to all those who always believed in our work and to all the INESC S3A/RVLab team who made it happen. Thank you Ana Isabel Torres, Ana Paula Costa, André Silva, António Cerveira Pinto, Alexandra Freitas, João Pereira, Joaquim Antunes, José António Conde, José António Proença, Manuel Eduardo dos Santos, Tiago Sepúlveda, and Vanda de Sousa, for making this IVE real. We hope everyone feels gratified.

Pathematic: a. rare (ad. GR.) pathematicos– liable to passion or emotions, s. pathema– what on suffer, suffering emotions, s. stem path– : see pathetic) pertaining to the passions or emotions; caused or characterized by emotions (*Oxford English Dictionary*).

Address correspondence to Carlos A. Martinho, Instituto de Engenharia de Sistemas e Computadores, Rua Alves Redol, 9, 1000 Lisboa, Portugal. E-mail: Carlos.Martinho@inesc.pt

project showed that the approach is a viable solution for short–time development of quality intelligent virtual environments.

For a while, Virtual Environments (VEs) were mainly centered around the interaction of single or multiple active participants with passive components of the VE. However, interest in user–independent components has grown to a point where it is almost impossible to consider building VEs without them. Synthetic actors are presently one of the most investigated independent components and are a valued feature in any VE. The following projects are but a few examples of VEs incorporating synthetic actors: the Jack system (Badler et al., 1997), SodaJack (Geib et al., 1994), Creatures (Grand et al., 1997), Virtual Theater (Hayes–Roth & van Gent, 1996), STEVE (Johnson et al., 1998), Herman the Bug (Lester & Stone, 1997) and Cosmo (Lester et al., 1997b), Edge of Intention (Loyall & Bates, 1993), ALIVE (Maes et al., 1997), Improv (Perlin, 1996), Oz (Reilly, 1996), artificial fishes (Terzopoulos et al., 1994), and artificial humans (Thalmann et al., 1997).

Naturally, developing intelligent agents for VEs raises a new set of development concerns. One of the major concerns involved in the creation of synthetic agents inhabiting virtual worlds is believability. Believable characters are characters that give the illusion of life and allow the user's suspension of disbelief (Bates, 1994). Participants are bound to virtual characters that have a consistent exteriorization of their personality, to characters displaying positive and negative reactions to the events of the synthetic environment that make sense when considering the perceived or assumed goals, beliefs, and attitudes of the characters (Ortony et al., 1998). To provide believability to the characters, techniques from different fields have and are being investigated. Drama and literary studies e.g., Hayes–Roth and van Gent (1996), artificial intelligence and artificial life techniques, e.g., Cañamero, 1997), and psychology models, e.g., Moffat (1997) and Reilly (1996) all contribute to the building of richer Intelligent VEs (IVEs).

Nevertheless, introducing higher level paradigms in the IVE also raises the performance overhead, a critical factor when developing IVEs for mid–range systems, or even when striving to penetrate the personal computer market. In these cases, architectures have to be designed around paradigms and metaphors allowing a reasonable performance achievement. However, the systems mentioned earlier were developed for high–end graphical systems and do not generally address such concerns.

Another important factor in the development of IVEs in the implementation time. The process is indeed complex, and building a graphical IVE from scratch is a task addressing several issues and problems from different fields of computer science. Strategies must be developed to allow the rapid development of such applications. However, most of the actual systems are built incrementally on previous systems, not always efficiently prepared for improvements.

This paper presents an approach based on techniques and models from artificial intelligence, drama, and psychology to create IVEs with believable emotional agents. The solution uses a theatrical metaphor instantiated over a modular architectural framework composed of three main modules: the dynamic scriptwriter (mind module), the theatrical company with its director and its cast of actors (body module), and the virtual stage managers (world module). The mind module is responsible for the narrative creation. The body module interprets it and performs upon it. The world module handles the camera and stage special effects. Furthermore, the communication paradigm of this architectural metaphor is aimed at the rapid development of real-time IVEs.

The paper focuses on the implementation of the mind module, which manages the emotional believability of the synthetic actors of the narrative and lays a set of concepts speeding the development process of these virtual agents. To achieve believability, the personality of the characters is expressed through the regular exteriorization of emotions reflecting the characters' inner goals, beliefs, and attitudes (Moffat, 1997). To support emotions, the mind module implements the theory of Ortony, Clore, and Collins (OCC) (Ortony et al., 1988) that handles the cognitive structure of emotions according to the inner character. The inner character implementation is based upon Frijda's two stage emotional stimuli appraisal theory (Frijda, 1987).

Finally, the paper describes the application of the architecture in the development of a real-time IVE. The application was developed in the context of the last world exposition of the century, Expo'98, and was implemented in less than four months, and exhibited on a midrange graphical system. Two synthetic dolphins, Tristão and Isolda, the main characters of the IVE, *lived* during the four months of the exposition in the virtual estuary of the Sado River and were displayed to more than one million visitors.

The paper is organized as follows. After a brief overview of the uses of emotions and personality in the field of believable agent systems, the scenario of the developed IVE is presented. Then, the architectural framework for short-time development of real-time IVEs is described as well as its instantiation for the Tristão and Isolda synthetic world. This description is followed by the explanation of the concepts underlying the implementation of the mind module and how it instills emotional believability in the synthetic actors of our virtual world. Finally, the results are presented and explained, and some conclusions are drawn.

RELATED WORK

The field of believable agents has suffered a notable growth over the past few years. When considering the development of IVEs with believable

agents, two main directions drive the investigators within the field. One direction follows a perspective based on biology, and more specifically, on neurology, to achieve rationality and believability in agent behavior. The other is a more empirical and artistic approach toward the agent-building process.

The first approach uses techniques and mechanisms from artificial life inspired by biological and neurological theories and postulates. One of the major sources of inspiration is the research of Damásio (1994). Along with several hypotheses in the field of neurology, it encourages the use of such mechanisms as neural nets and reinforcement learning to implement the agent behavior model. In these IVEs, the agents develop, on their own, strategies to cope with the characteristics of the synthetic world they inhabit and achieve their inner goals. The process is done without user supervision. Additionally, the information gathered throughout the observation of one agent generation can be reused to create a new set of better adapted newborn agents using mechanisms such as genetic algorithms. The attractiveness of this approach resides in the automation and independence of the development process from direct user supervision. One can just sit back and watch the agents (eventually) develop interesting and rational behavior. Examples using these approaches can be found in Gridland (Cañamero, 1997) and the Cyberlife Technology commercial product, Creatures (Grand et al., 1997).

However, when the development goal is to build agents with a specific and predefined personality, this approach is not the most adequate. To reach the desired personality, the agents would have to be submitted to a set of empirical test environments, thoroughly designed and planned for the purpose. The creation of these situations is hardly practical: they are not easily controllable and do not even ensure the results.

Thus, there is a need for a more practical and immediate approach, an approach based on empirical evidence. Following an artistic perspective, a second approach has appeared. Like actors who must interpret a well-defined role in a play, this direction of investigation focuses on the elaboration of strategies allowing the configuration of the behavior of the synthetic characters to achieve a consistent exteriorization of their personae. Often supported by theatrical metaphors, this second line of investigation also takes theories and postulates from psychology to model the expected personality of the synthetic actors. Although very practical and artificial, this approach is often necessary, mainly when developing agent-based applications, where the agent function is critical and its behavior must be predictable and consistent.

A set of strategies has been developed to create specific personalities for synthetic actors. At Carnegie Mellon University, Bates' group started using linear variables to define the mood of *woggles*, agents inhabiting the Edge of Intention (Loyall & Bates, 1993) virtual world. The variables were used to

represent the physical, social, and emotional dimensions of the synthetic characters. At Stanford's Knowledge System Laboratories, the team of Hayes-Roth took the idea one step further, incorporating linear personality traits to define agent behavior in the Cybercafé (Rousseau & Moulin, 1996; Rousseau & Hayes-Roth, 1996). Although the personality trait models have a great appeal in psychology, they remain too generic a way to define one-character personality. Another direction was undertaken with Reilly's work (Reilly, 1996) at Carnegie Mellon University where, instead of defining personality, the emotion generation process was implemented following the OCC theory of emotions (Ortony et al., 1988). The result, the Em system, was implemented atop the Tok architecture (Loyall & Bates, 1993) and allowed the fusion of emotions in the agent design. Another important piece of work in the field is the Will architecture (Moffat, 1997). With Will, Moffat followed a similar path to Hayes-Roth's Virtual Theatre (Hayes-Roth & Rousseau, 1996), also using personality traits to define the personality of the agents in very specific situations. However, in this particular case, the Big-Five psychology personality model was chosen instead of empirical variables adapted to the specific implemented situation. Two other features of this work are also important. First, a two-stage emotional stimuli appraisal based on Frijda's (1987) theory of emotions was used to develop emotions beyond the reaction process. Second, emotions and personality are seen as two interrelated concepts. Personality and emotion are two views of the same concept differentiated by two dimensions: duration and focus (Moffat, 1997).

If most of the research undertaken until now falls in the entertainment field of application, one must not underestimate the applicability of believable agents. The Microsoft Persona (Ball et al., 1997) project is a good example of the potential in using emotional and personality approaches in building interface agents. The field of Intelligent Tutoring Systems (ITS) has gained much in using believable tutor agents. Some analysis of the Persona effect has been performed upon ITS improvement (Lester et al., 1997a). Herman the Bug (Lester & Stone, 1997) and Cosmo (Lester et al., 1997b), STEVE (Elliot et al., 1997), and Vincent (Paiva & Machado, 1998) are good examples of using personality to convey tutoring agents' believability and producing better tutoring systems.

Believable agents may well be the key for information cost structure reduction, the motivating force of 21st century graphical user interfaces. In providing a higher abstraction level over information, they will allow better perceptual control over its complexity, enabling the manipulation of new thought structures.

SCENARIO

Lisbon held the last world exposition of the 20th century: Expo'98. Its theme was "The Oceans, a Heritage for the Future." One of the greatest

symbols of the oceans is the dolphin. However, dolphins could not be brought to Expo '98: they are far too sensitive mammals to be submitted to the stress of the exposition without risking harm. Synthetic dolphins were the only chance for dolphins to be represented in the exposition. The Territory pavilion of Portuguese representation wanted not only to fulfill the exposition theme, but also to capture pavillion visitors' interest in Portuguese scientific projects. The conjunction of these factors launched the development of S3A, an interactive application featuring artificial dolphins.

A team from heterogeneous fields of expertise was assembled. Biologists and ethologists from the Delphim project, a nonprofit-making organization devoted to the study of the Sado dolphin community, assisted the engineers in the simulation of realistic behavior and motion for the dolphins. Designers and artists from the Territory pavillion assisted in the creation of interesting personae for the characters, in the building of an involving background history for the characters and in all the artistic aspects of the system. All gave their contribution to make this project believable and interesting, a good balance between artistic and realistic virtual expression. Tristão and Isolda were born.

The application S3A takes us to the beginning of Ages, to a place where humans and dolphins had learned to communicate with each other— Atlantis. As the civilization flourished, the Atlanteans developed an apparatus that enabled them to express their feelings to dolphins. Humans and dolphins built a relation of friendship and trust that endured till the civilization disappeared. Fragments of the Atlantean civilization were gathered throughout the ages in Portugal and the fragment analysis enabled the construction of human-dolphin communication apparatus. S3A enables the experience of communicating with dolphins. It allows the user to influence the behavior of a bottlenose dolphin whose descendance goes back to the dolphins that once populated Atlantis waters— Isolda.

Isolda is a common Atlantic dolphin, quick and agile, who arrived on the Portuguese shore where she met Tristão, a male Sado bottlenose dolphin. Isolda, due to her descendance, has faith in mankind and still has embedded in her genes the ability to sense the action of the apparatus. Tristão is a typical Sado dolphin— heavy, shy, and very suspicious of man's activities. Fishermen killed many of his community throughout this century. He is insensible to the apparatus.

Isolda, Tristão, and the human presence (via the apparatus) are the vertexes of a dramatic triangle introduced to involve the public in S3A's plot. Tristão and Isolda are represented in Figure 1.

The communication apparatus, a ceramic sculpture of a dolphin equipped with four pressure-button sensors, was laid in the middle of the exhibition room, in front of the screen as shown in Figure 2. The visitor can, at any time, approach the sculpture and touch one of the buttons to influ-

FIGURE 1. Tristão (left) and Isolda (right).

FIGURE 2. Interface.

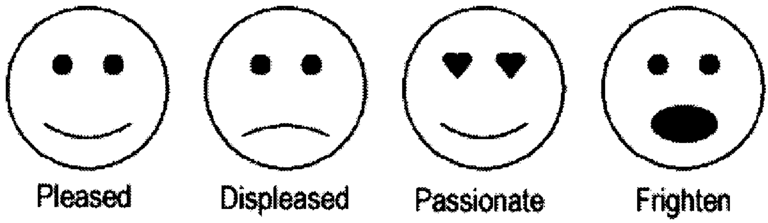

FIGURE 3. Sensors.

ence Isolda's emotions. The four types of emotions that the visitors can express to Isolda are represented in the four sensors of Figure 3.

ARCHITECTURE

Developing complete graphical IVEs with believable agents from scratch is a complex process, addressing several issues and distinct fields of computer science. The architecture underneath the IVE must support subdivisions centered on field-specific premises to be tackled by field-specific development teams. Furthermore, the architecture modularization must ease its parallel implementation, thereby reducing critical paths in the development process and contributing to a shorter development time. Strong module independence also reduces communication problems between teams and allows each development nucleus to explore adequate strategies to tackle their specific problems.

Another important issue in IVE development, mainly when using low- to mid-range systems, is performance. Since most of the processing time will be spent with the unavoidable graphical calculations and rendering process, the architecture must support a lightweight agent implementation paradigm.

The architecture presented in this section addresses these issues and is therefore aimed at:

- short development time;
- efficiency and real-time achievement.

To achieve these goals, the architecture is based on a theatrical metaphor striving for a concurrent modular development process.

Theatrical Metaphor Overview

The theatrical metaphor is instanced over a modular architectural framework composed of three functional units (modules): the dynamic scriptwriter (mind module), the theatrical company with its director and its cast of actors (body module), and the virtual stage managers (world module).

Figure 4 represents the overall architecture. The designations *mind*, *body*, and *world* are used to provide a bridge to the vocabulary used in most agents systems, but their meaning is slightly different in the context of this paper:

- The *mind module* is responsible for the narrative creation. It manages all the agents at the narrative level and controls the emotional believability of the characters.
- The *body module* interprets the narrative and acts on it. It manages the geometrical and audio–visual planning and controls the life–like believ–ability of all the characters.
- The *world module* handles the virtual cameras and the stage special effects. It manages all aspects related to audio–visual display of the character per–formance.

Each module is described in more detail in the next subsections.

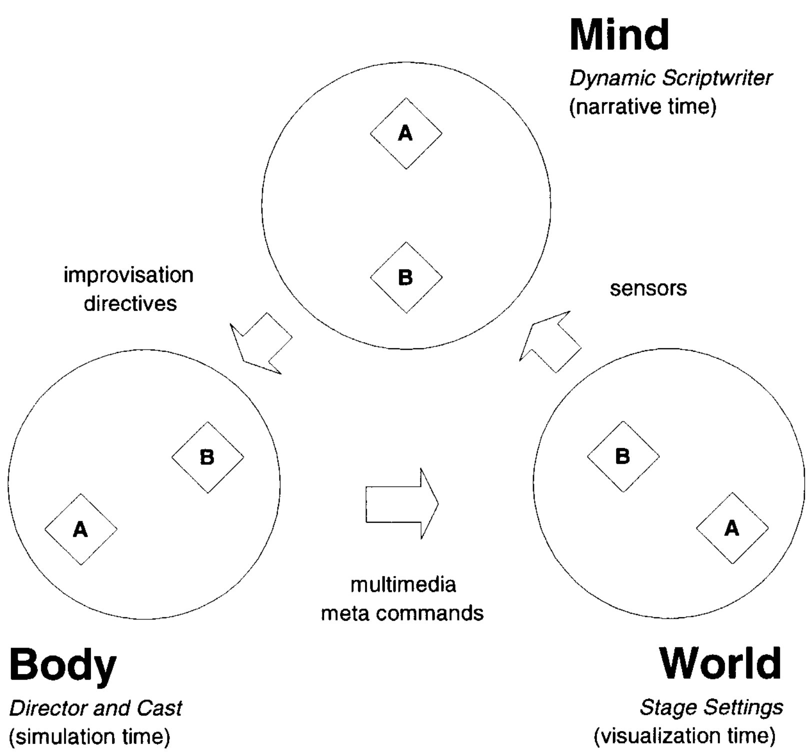

FIGURE 4. Architecture overview.

Communication

Each agent is implemented by three parts: a mind part, a body part, and a world part. Each part is a specific image of the agent, and each image is handled by one of the previously described modules. Therefore, each module implements a specific part of all the agents of the IVE.

The architecture allows a one-way circular communication channel between the modular images of a single agent, as can be seen in Figure 5. This one-way communication paradigm reduces development and implementation synchronization problems. Naturally, this has a cost. In this case, the cost is information redundancy to preserve the context in each module. However, considering that the reduction in communication and synchronization between modules also benefits the system performance, this is an acceptable drawback. Furthermore, the architecture allows every agent image to communicate with all the images of the same type of the other agents (implemented by the same module). The communication paradigm between agent images is also depicted in Figure 5.

This second communication aspect is based on information proximity. The assumption is that each image is interested in a predetermined type of data likely to interest only parts of the same type. This approach allows the reduction of communication between modules and the tight integration of

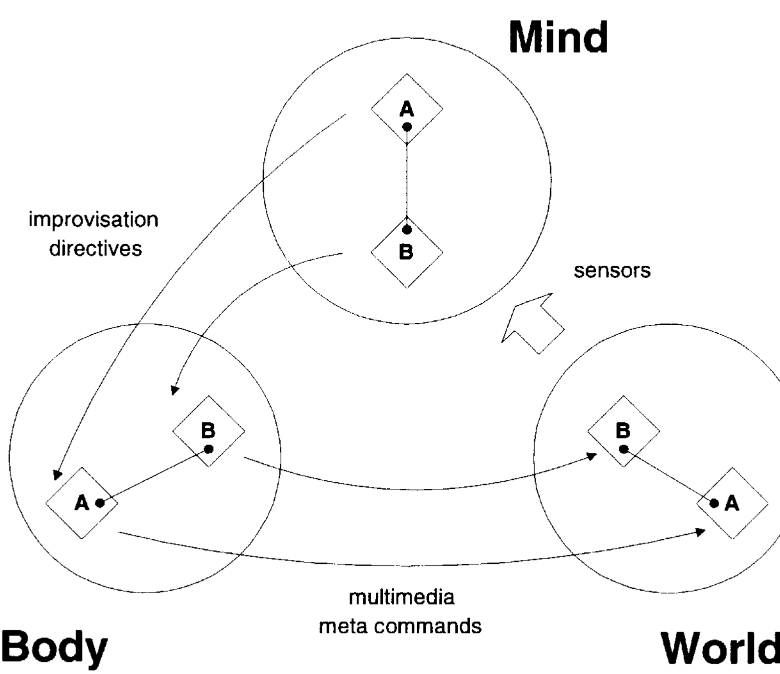

FIGURE 5. Module communication.

information in a highly efficient implementation, independent of any other module implementation.

Communication is an important issue when considering the short development time of a real-time IVE. This communication paradigm allows each module to efficiently manage the manipulated data type without worrying about the other modules.

Mind Module

The mind module implements a dynamic scriptwriter. As a scriptwriter, it is responsible for the generation of the storyboard. The storyboard is implemented as a set of improvisation directives for all the characters of the narrative world, much like a scriptwriter writes the storyboard for a play. These improvisation directives are the data passed on to the body module for actual performance. This is done dynamically; each time the scriptwriter creates new improvisation directives, it takes into account all sensor data received since the last generated set.

The time span between two consecutive steps of the mind module is designated the *narrative time*. It is a subjective time span that represents the time for something relevant for the storyboard to be planned and executed. This time span is irregular. For instance, consider the following two examples of improvisation directives.

Example 1:

```
⟨Isolda swims very slowly to the plane wing⟩
⟨Tristao slows down to Isolda speed⟩
⟨Tristao follows Isolda closely⟩
```

Example 2:

```
⟨Isolda swims near the bottom⟩
⟨Isolda swims to rock⟩
⟨Tristao stops swimming⟩
⟨Tristao watches Isolda⟩
⟨Tristao move its dorsal fins nervously⟩
```

It is not obvious which set of improvisation directives will take more time to be fulfilled. The distance to the plane wing, to the rock, and the speed of the dolphins, are all factors influencing the time of completion. As opposed to literature, where narrative-irrelevant and long actions (in terms of the real time it would take to perform them) are summarized in a few words, the IVE must perform them in real-time. As the details of the performance of these directives are irrelevant to the narrative, the dynamic

scriptwriter will simply keep on storing sensor readings until the actual improvisation directives are performed and a new set must be generated. In other words, the mind module has a lightweight implementation and is not processed at each simulation cycle. This feature allows its integration in a real-time graphical intensive IVE such as Tristão and Isolda.

Body Module

The body module implements a theatrical company director and her set of actors. The improvisation directives received from the module are the set of commands that the actors must fulfill, much like real world actors follow their script. They must follow it accurately but can improvise on it as long as they still convey to the viewer the actions specified by the improvisation directives and stay "in character." This is called the *directed improvisation paradigm*. The classical Commedia del'Arte established a set of improvisation guidelines written to help actors in achieving these goals, which are applied when implementing the body module. Like their human counterparts, the agents must work together to produce engaging performances. They must accept all offers, they must not block other actors, they must act naturally, they must not try to be clever, and they must try to reincorporate previously generated elements. These improvisation guidelines have also been successfully used in the Virtual Theatre project (Hayes-Roth et al., 1996).

To help achieve acceptable improvisational character interactions, the concept of a *director* was used. The director is responsible for handling each actor's performance in such a way that each character's sequence of actions does not hinder other characters' performances. Thus, the director resolves conflicts between characters and directs the character actuation. Note that it is possible to implement the body module without a central director.

The body module ends up controlling the simulation time. It computes audio and geometrical primitives and passes them to the world module. It also monitors the progress of all improvisation directives. When all have been performed, the body module director fetches a new set of directives from the mind module and activates the actor planners with the parsed improvisation directives and further performance direction. Since new improvisation directives are only fetched when all actors have performed their actions, some actors will perform quicker than others. At this point, improvisation comes into play and the guidelines must be followed. Actors must continue to perform according to their recent actions and emotional state, until a new set of primitives is retrieved from the mind module.

The time span between two consecutive steps of the body module is designated *simulation time*. It is the time used in each simulation cycle for actuation planning and geometry and audio calculations.

World Module

The world module implements the stage management and is responsible for translating the low-level meta-primitives sent by the body module into audio-visual output. It is also responsible for reading the sensors, data from which is passed at each cycle to the mind module, and used to handle other stage-related visualization aspects.

Several abstractions fall under the world module responsibilities. Handling the camera according to what is going on (helped by some external direction) and producing special effects to enhance realism are two examples of typical world module tasks. For instance, in Tristão and Isolda, the camera follows Tristão and Isolda and tries to keep them in its field of view– they are, after all, the main characters of the IVE. When it is not possible, it focuses on Isolda, who is the target of sensor actuation. The camera may also receive commands from the mind module (via the body interpretation) to follow a specific character when it is critical for the narrative perception, like following a fish when it is about to be eaten. Additionally, the world module also manipulates a set of fish schools and underwater creatures to which characters are completely indifferent. Like the lighting effects, they exist only because of aesthetic reasons, not narrative ones.

The time span between two consecutive steps of the world module is designated *visualization time*. It is the time used in each simulation cycle for geometry rendering, audio playing, and sensor reading.

MIND MODULE IMPLEMENTATION

The mind module implements the dynamic scriptwriter and is responsible for the believability of the narrative characters. It handles all the decisions related to narrative-relevant aspects of the IVE, generating the storyboard for each character according to their defined persona. To achieve this goal, the mind module manipulates a set of concepts that abstractly represent what is happening in the virtual world. The principal concepts are explained in the following subsections. More on the concepts and the mind module implementation can be found in Martinho (1999).

Concepts

World Abstraction

The world is modeled by a set of concepts representing the topology of the IVE. It is composed of *areas*, which are organized as graphs where each node represents an abstract location and each arc represents accessibility

between the linked nodes. Areas also define visibility. A character can only see the events and entities of the location at which it stands. The definition of areas is more than a simple geometrical topology mapping.

For instance, the virtual estuary of the Sado River is implemented by a grid of areas, although no physical obstacle, except the plane pieces and some rocks on the bottom, really block the character vision. However, the level of pollution of the waters appreciably decreases the vision distance and is implemented by distinct areas. On the other hand, the dolphins' sonar system allows them to overcome this difficulty, so some actions are implemented without visibility restrictions. Fishes, however, have a hard time seeing predators come. The virtual world of S3A is represented in Figure 6.

Several entities coexist in these areas. *Entities* are:

- the *characters*, the synthetic agents inhabiting the IVE;
- the *décor*, the relevant passive elements of the IVE.

Thus, Tristão, Isolda, and their prey are the characters of S3A, while the rocks and the plane crash pieces are the *décor* of S3A. Some décor elements may be manipulated and/or may have containing capacity.

Entities are organized as a class hierarchy. This classification allows reference to groups or group elements when manipulating information in the IVE. For instance, the entity hierarchy represented in Figure 7,[1] a simplified set of the underwater taxonomy used in S3A, allows the manipulation and

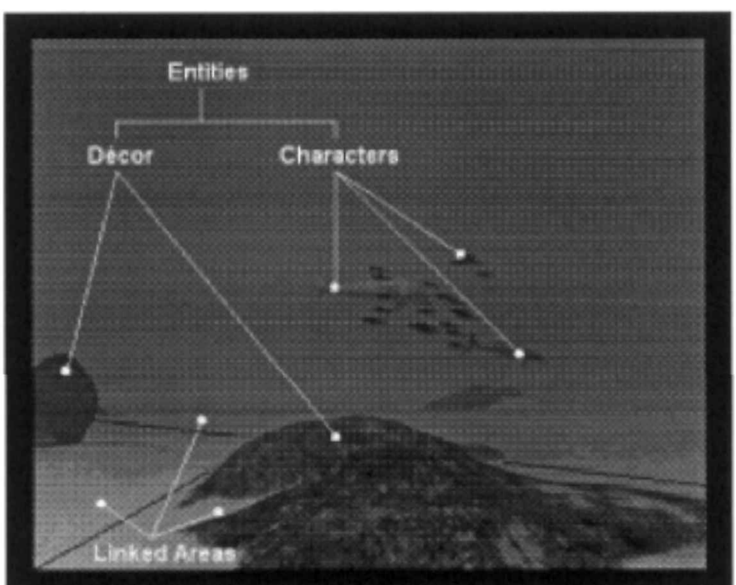

FIGURE 6. S3A world abstraction.

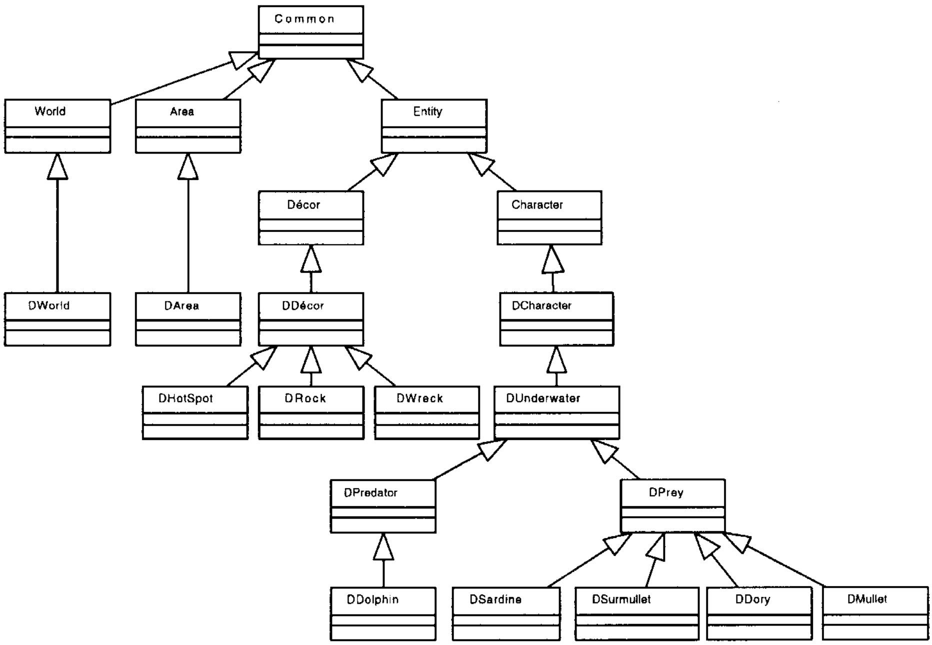

FIGURE 7. World implementation.

expression of concepts predator and prey. The hierarchy defines the concepts that can be used for knowledge manipulation. The right concept structure is fundamental in achieving the pretended IVE behavior.

Actions and Events

Everything happening in the IVE is an *event*. Events are simply the actions of the entities of the IVE. They are a perceptual abstraction. In our underwater world, a simple event could be:

```
([gen77]   of  EVENT
 (time         77)
 (subject      Dtristao)
 (id           Afollow)
 (objects      Disolda)
 (parameters  veryClose)
 (location    LsouthEastBottom)).
```

This would represent that Tristão, our male dolphin, is following Isolda, its female counterpart, very closely near the bottom of the sea in the Southeast area of the IVE.

At the end of the narrative cycle all the events are converted to improvisation directives. For instance, the previous event would be translated into.[2]

```
[Tristao] follow [Isolda] neardistance-3 [BOT_SE]
```

and then sent to the body module.

Personality and Emotions

The characters of the IVE must be believable. To achieve believability, one must define the persona of each character and ensure that this persona is correctly exteriorized. However, defining a character personality can be viewed from another perspective– personality as the regular or expected expression of emotions. As stated in Moffat (1997), personality and emotions are very close although they appear very different.

- Personality is defined in terms of temporal consistency, but emotions are defined in terms of temporal inconsistency.
- Emotions are focused on a particular event or object, while personality can only be described in general terms.

Assuming that the underlying mechanisms are both essentially the same, one can unify the phenomena and consider *duration* and focus to be two dimensions plausibly identifiable at the cognitive level. Following this idea, personality emerges from emotion consistency and can be expressed by a set of typical reactions to specific events. As opposed to personality theories, emotion theories are more generic and implementation-friendly. OCC theory provides a theoretic framework for handling emotions and their inner interactions with goals in terms of behavior repercussion.

In S3A, the main characters are defined using the Myers-Briggs[3] and Big-Five[4] theories of personality, instanced to specific situations where the personality comes out clearly. Then, they are converted to typical emotional reactions and their respective exteriorization to be implemented by the IVE, according to the OCC model. This approach is similar to the ones used by animators, who define characters using a typical set of situations and reactions.

Character Cycle

As with the character persona implementation, the character behavior implementation is inspired by a human psychology model. The implementation is based on Frijda's theory of emotions (1987), which provides the basis for the character perceive-react-reason-and-act cycle. This psychological theory of emotions complements the OCC model that handles the emotional structure.

Each character cycle is thus divided into four phases:

1. In the *perception phase*, the events of the world are filtered according to the characters' interests and location. Each character also updates its model of the world at this stage according to its immediate perception.
2. In the *reaction phase*, the immediate impact of the events on the characters is estimated and emotional responses are appraised.
3. In the *reasoning phase*, the implications of changes in the internal world model and on every aspect of the character life are evaluated.
4. In the *action phase*, the character finally acts based on the evaluation performed in the former steps.

The four phases are explained in more depth in the next sections.

Perception Phase

Character perception is based on two principles: characters do not perceive all the events; characters do not perceive the virtual world as it is.

Limited Perception

The characters are assumed to have a limited perception capacity. They only perceive events happening in the area where they stand. Although this approach does not allow a distinction between different character perception capacities, it implements the notion of a *theatrical scene*– a virtual space, where all characters share the events, but where they are also indifferent to events happening in other scenes. Of course, this approach can be overridden at implementation. For instance, the implementation of the dolphin sonar capacity overrides the notion of scene area. Additionally, characters only perceive events that are relevant for their concerns (Frijda, 1987). The events of the world are chosen through a set of *filters*, which are specified as a set of action patterns and represent the character interests, major needs, and preferences.

Inaccurate Perception

Each character has *entity models*, which are subsets of parameters associated with an entity that are updated when the entity is perceived. They represent the immediate aspects of entities relevant to a specific character and allow characters to perceive an entity as being different from how it really is. Each character selectively divides the perceived events per entity. For each entity all related events– that is, all events performed by the entity or for which the entity is the object of the action– are mapped to world model attributes. Color is an example of an immediate aspect of an entity. The nonimmediate aspects of the entities are handled later, in the reasoning phase. The final actuation decisions, performed in the action stage, are always based on the world models, never directly on the events perceived from the IVE.

Implementation

The perception phase is implemented by three sequential steps as follows:

1. *Visibility filtering*: all events and entities falling outside the character area are discarded.
2. *Interest filtering*: all events not matching the character filters are discarded.
3. *World model update*: the model for each perceived entity is updated.

Figure 8 summarizes the perception phase.

Reaction Phase

Character reaction is based on three principles: characters react based on inner goal repercussion; characters react emotionally to their surroundings; each character expresses its emotions differently.

Inner Goal Repercussion

Characters have *goals* whose implementation is based on OCC theory. Each goal is implemented with three important attributes: the *likelihood of success* quantifies the perceived chances of the goal succeeding; the *importance of success* quantifies the relevance for the character of being well succeeded in achieving that goal; and the *importance of failure* quantifies the relevance for the character of failing in achieving that goal.

Although the likelihood of success and the likelihood of failure are complementary concepts, importance of failure is not a complementary concept to importance of success. For instance, dolphins have a replenishment goal for breathing that has a cyclic activation period. As with humans, it is natural for dolphins to breathe. Thus, success here is not important for dolphins: they will not feel joy each time they breathe. However, if a dolphin

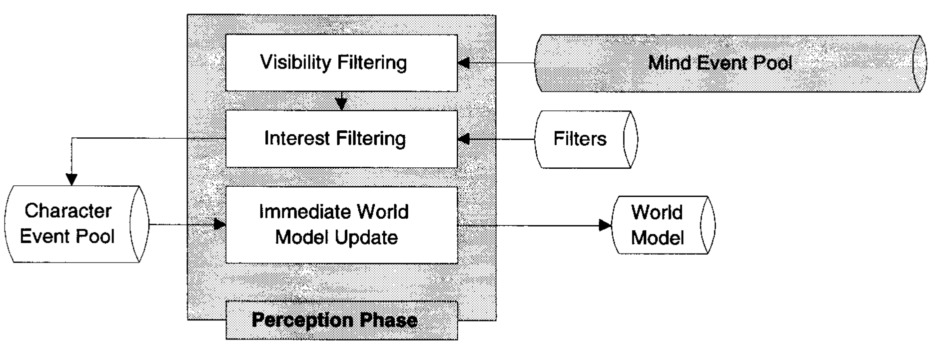

FIGURE 8. Perception phase.

perceives itself as unable to breathe, it will most certainly feel an intense fear, not to say panic. Failure in breathing is an important matter and managing to breathe in these situations will provoke very strong relief.

Goals are organized as a graph, as OCC theory suggests. The goal/subgoal relationship is implemented with two goal lists for each goal. Generally identical, they represent the goals activated on the goal activation and the goals deactivated when the goal is deactivated. But goals are also connected through four types of semantic links.

- *Necessary link*: if the goal fails, the linked goal also fails.
- *Sufficient link*: if the goal succeeds, the linked goal also succeeds.
- *Facilitative link*: if the goal succeeds, the likelihood of success of the linked goals increases by a predefined amount.
- *Inhibitory link*: if the goal succeeds, the likelihood of success of the linked goals decreases by a predefined amount.

When a goal succeeds or fails, the goal graph is updated according to the semantics of its links. For instance, the *breathe* goal has a subgoal: *move-up*. If *breathe* is activated, so will *move-up*, the same happening on deactivation. *Move-up* is a necessary goal for *breathe*: if *move-up* fails, so will *breathe*.

Emotional Reaction

Characters display *emotional responses*. The mind module implementation considers the 22 base emotions of OCC cognitive theory. As proposed in this theory, emotional reactions are classified into three categories: emotions based on aspects of objects, appraising the liking of entities with respect to the agent's attitudes; emotions based on actions of agents, appraising the approval of the actions with respect to the agent's standards of behavior; and emotions based on consequences of events, appraising the pleasingness of events with respect to the agents goals. The complete classification is presented in Figure 9. The emotion intensity depends on the events, goals, and entities responsible for the emotion.

Each character has a set of emotion thresholds and decays according to its personality, defining its emotional profile. The threshold values represent the character resistance toward an emotion type. It is only when the intensity is above that threshold that the emotion is active. Decay represents the emotional memory of the character for the emotion type. The higher the decay rate the quicker the character will "forget" the emotion.

The creation of emotions is handled by the emotional reaction mechanism. Emotional reactions are an automated mechanism following OCC theory that launches emotions according to the events perceived and the emotional characterization of the actors. The four emotional reaction types are listed below.

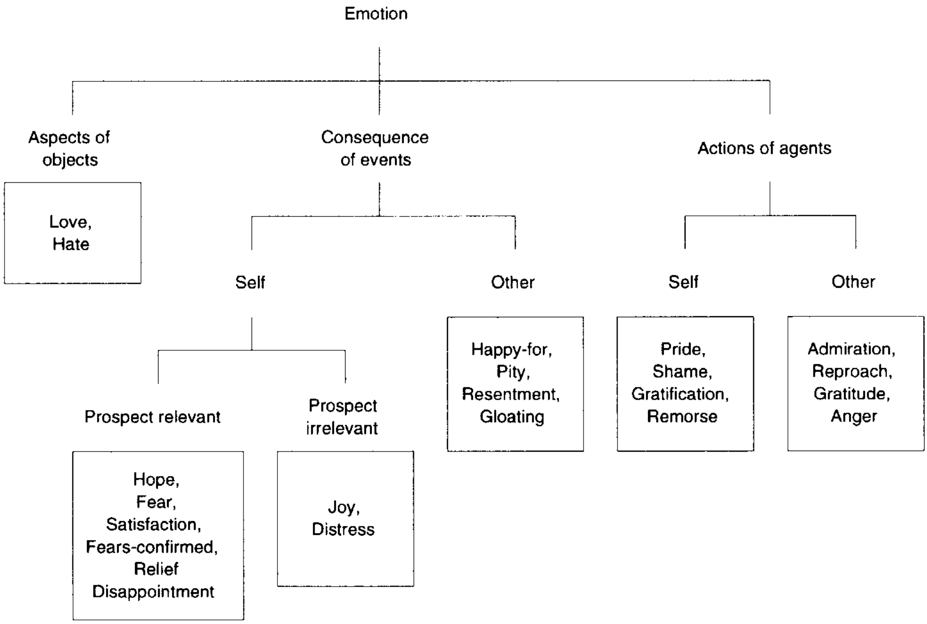

FIGURE 9. Emotions by type.

- *Attraction reactions* handle emotions related to aspects of objects. They are triggered by the proximity of relevant entities.
- *Event reactions* handle emotions related to consequences of prospect-irrelevant events. They are triggered when a perceived event matches a predefined pattern.
- *Prospect-based reactions* handle emotions related to consequences of prospect-relevant events. They are triggered by goal changes.
- *Attribution reactions* handle emotions related by actions of agents. They are triggered when a perceived event matches one of the character standards of behavior.

The emotional character's profile is mapped into a set of reactions triggered by entity proximity and event patterns. Again, however, emotions are not only a reaction process. Emotions resulting from thought, as, for instance, hating someone for sending someone else to do bad things for her, must be implemented later in the cycle, in the reasoning stage.

The emotional reaction process is illustrated in Appendix A. Each example is presented in two parts. First, a textual explanation from the narrative point of view is presented. Then, the reaction process is summarized in a table. The tables present, from an implementation point of view, the three elements responsible for the launch of emotion: the reaction involved, the

event matched, and the emotion launched. Note that the implementation is simplified for the sake of clarity. More details of the implementation can be found in Martinho (1999).

Individual Emotional Expression

The concept of *behavior features* enables different exteriorization of the same emotion by different characters, even when under the influence of the same emotions with the same intensities. Behavior features were suggested in the OCC theory and are inspired from the implementation in Reilly (1996). They are an indirection layer, mapping the emotions into a set of concepts representing aspects of the personality of each character.

Behavior features are calculated based on the intensities of active emotions. Each behavior feature is the weighted sum of all emotion classes. Naturally, the organization by type of emotion is not adequate for behavior feature implementation. Figure 10 depicts the classification of emotions by valence, which is more appropriate for behavior feature handling.

The dolphins have four behavior features mapping each sensor:

- Pleased: the sum of all positively valenced emotions.
- Displeased: the sum of all negatively valenced emotions.

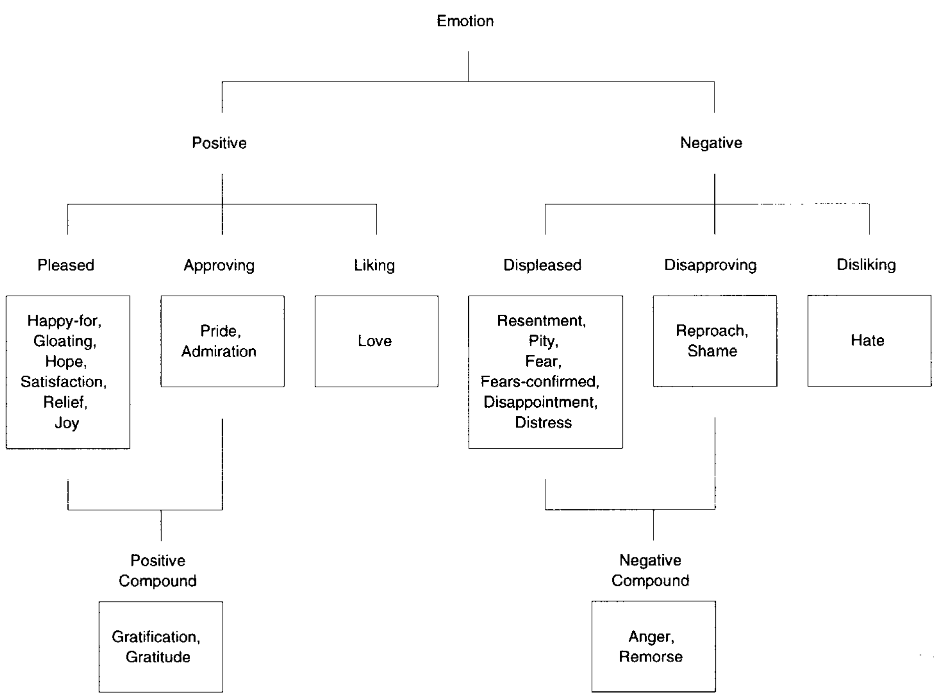

FIGURE 10. Emotions by valence.

- Passionate: the sum of love emotions.
- Frighten: the sum of fear emotions.

Actuation decisions are always based on behavior features, never on the emotions themselves.

Implementation

The reaction phase is thus implemented according to the following steps, illustrated by Figure 11.

1. *Goal update*: all character goals are checked for success or failure deactivation and change of likelihood according to the events and the world model. Each detected change produces a new event to be handled in the next step.
2. *Reaction check*: the characters react instinctively to the set of perceived events and goal alterations. All reaction type are checked for emotion launch according to the character's emotional profile.

Reasoning Phase

The reasoning process is performed by a set of production rules whose preconditions are based upon the world model, the goal state, the behavior

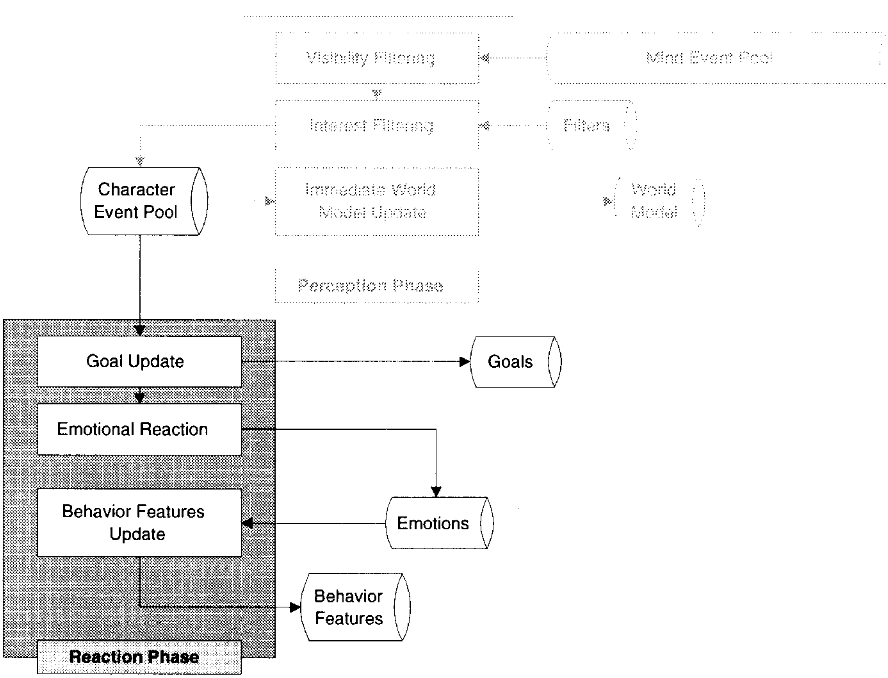

FIGURE 11. Reaction phase.

features, and specific inner state data. The results of the reasoning process update the inner character state and its world model. The emotional conditioning of the characters is implemented by changing the emotion thresholds and decay rate to reflect a particular mood. New emotions can be generated at this stage, but the behavior features are not updated to reflect these new emotions in the same cycle. Figure 12 shows the reasoning phase.

Action Phase

Character action is based on two principles: characters act according to resource availability and characters act according to action importance.

Resource Availability

Characters use *resources*, an abstraction over the availability of means to perform an action. In each narrative cycle, only one action can use each resource. Therefore, characters can perform several actions in the same cycle

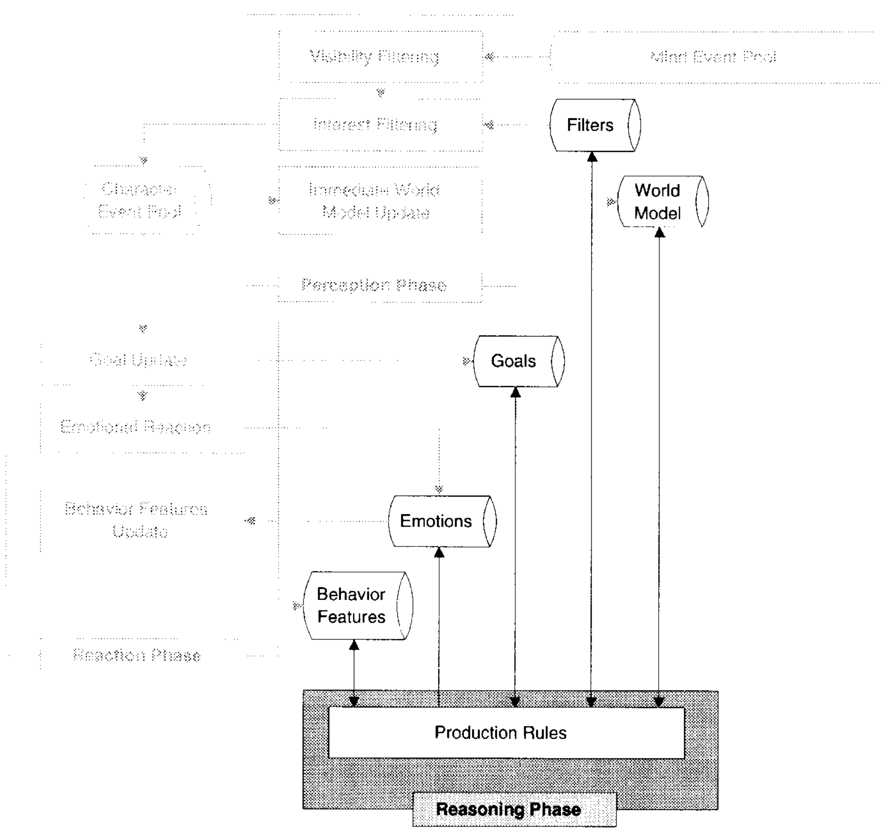

FIGURE 12. Reasoning phase.

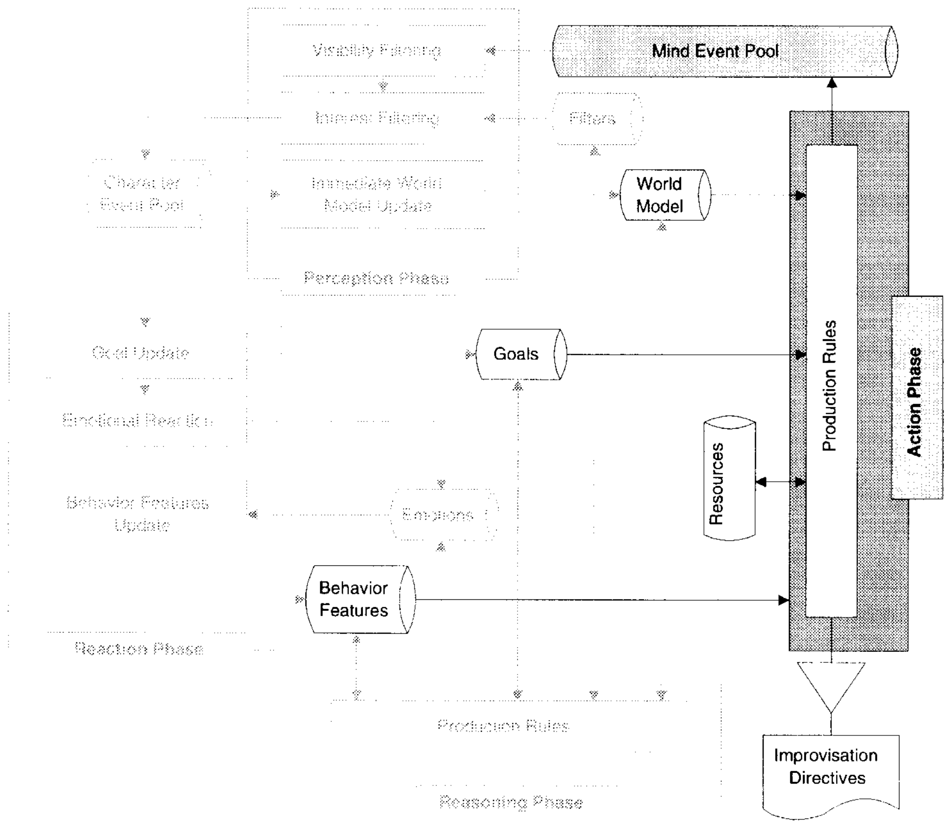

FIGURE 13. Action phase.

only if the actions do not share the same resource. This approach is inspired by Alive's Silas T. Dog (Blumberg & Galyean, 1995) and the Improv System (Perlin, 1996).

In S3A, the dolphins have a resource for each part of their body: the head, the mouth, each one of the fins, and the tail are the dolphin's physical resources. If the mouth is used to eat something, then it cannot be used for anything else in the same narrative cycle. Note that resources can also be used to represent other exclusive decision mechanisms and not only physical restrictions.

Action Importance

Each goal has a *priority*. Priority reflects the importance of the goal for the character, the preference amongst goals active at the same time. The approach of actions competing between themselves for execution and resources is inspired in Brooks' subsumption architecture (Brooks, 1986) and Blumberg's multilevel goal architecture (Blumberg & Galyean, 1995).

Implementation

As reasoning, action planning is implemented by a set of priority production rules, reflecting the goal structure of the agents. Each goal has a set of exclusive production rules reflecting the several aspects of the goal implementation. When a goal becomes active, the associated production rules become active. Production rules can have behavior feature restrictions. In this case, all rules not verifying them are deactivated. Active production rules are organized by resource. For each resource, the rule for the highest priority goal is fired– the character acts. Note that goal activation and deactivation is not handled by production rules but in the reaction phase. Finally, actions are exteriorized as a set of improvisation directives sent to the body module, but also to the mind pool of events so they can be used as sensorial information in the next character cycle. Figure 13 illustrates the action phase.

RESULTS

S3A ran nonstop during the four months of Expo'98 and was displayed for an average of 10,000 visitors per day. Two aspects must be considered and evaluated separately:

- the technical viability of the architecture and the mind implementation
- the emotional believability of the created characters.

Technical Viability

Development

S3A was developed in less than four months by a team of 15 professionals working part-time on the project. The team was composed of:

- 1 project manager
- 3 ethologists
  - 2 dolphin geometry and behavior consultants;
  - 1 dolphin sound consultant.
- 3 artists
  - 1 designer, responsible for all geometry and special effects;
  - 1 storyboard writer, responsible for the creation of the character personae and for the possible scenario developments;
  - 1 architect, responsible for the overall design of the room where the system was installed.
- 7 software engineers
  - 1 mind module developer, responsible for the character narrative personae implementation;

  - 3 body module developers:
    - * 2 developers for the underwater creatures' physical movements and responsible for the life-like quality of the dolphins;
    - * 1 developer for the planners, responsible for the mapping of improvisation directives into "real-life" movement primitives.
  - 3 world module developers:
    - * 2 developers for the geometrical integration and visualization as well as for the intelligent camera development;
    - * 1 developer for the 3D sound server.
- 1 electronic engineer, responsible for sensor integration.

The module independence allowed the parallel development of the three modules, thus reducing the overall development time. Table 1 shows the effort spent in each of the development areas.

As expected, most of the effort (more than 50% of the project effort) was spent on geometrical issues of the implementation. Most of the personality implementation time (60% of the mind module development time) was spent on parameterizations, tests, and refinements (e.g., the OCC variables and emotional reaction parameters).

Since the project weight was on the visualization aspects, the insertion of other agents with already implemented bodies, even with very distinct personalities, would involve a relatively small effort. Only performance limitations prevented the "birth" of a third dolphin– Junior.

Performance

The application ran on an Intergraph Realizm2 at 1024×768 true colors. The final frame rate oscillated between 9 and 14 frames per second, with an average of 12 frames per seconds. The final scene complexity was 25,000 polygons with 40MB of textures. The application load was divided over the two processors. The first worked full-time to cope with the physical

TABLE 1 Development Effort

Task	Person-Month
Management	1
Dolphin consulting	0.5
3D design and FX	0.625
Storyboard creation	0.25
Mind module	2
Body module: physical movement	8
Body module: planning	2.5
World module: 3D sound	1
World module: camera and integration	4
Sensors	0.125
Total (without installation)	20

calculations of the two 4,000 polygon dolphins and the average ten 600 polygon fish characters. The second processor, dedicated to all the other implemented tasks, worked only at 80% on average. Also, as expected, the bottleneck was on the physical calculations, and the mind module did not introduce any performance degradation (more information on the physical model implementation can be found in Sepúlveda 1996).

The mind implementation confirmed the viability of integration in real-time 3D IVE with minimum performance degradation on a midrange graphical system. The main benefit of using this architecture was that the one-way localized communications between modules allowed great independence over implementation and allowed liberty in achieving efficient algorithms and data representation for each module purpose.

The principal flaw of this approach was the need for information redundancy– the body module needed to store some context information already stored by the mind module. However, the time saved in protocol development and synchronization implementation compensated for this, without even considering the possible serialization side effects. The freedom of efficient development of strategies, not attending too much to other modules, was critical to the short development time and real-time achievement.

Emotional Believability

The prime requisite of the project was that the simulation had to be faithful to real life. Dolphins could not perform actions that they would not do in the real world. Emotional expression could not be exaggerated to convey clearly the emotional state of the characters, an effective technique used by character animators (Bates, 1994). The challenge was to build on the real world behavior to make the visitor understand the human personalities underneath the dolphin skins, conveyed by real life emotional exteriorisations.

Before the installation at Expo'98, the agents' behavior was statistically confirmed. The simulation analyzed two factors (goal selection and behavior feature intensities) over 400,000 narrative cycles (approximately four months of simulation time), to get a feeling for the would-be behavior of the characters of S3A during Expo'98. This (Martinho, 1999), along with some preliminary results obtained during the installation of the system, were a good "guarantee"[5] that the dolphins would behave according to their predefined personalities.

After an initial two-week evaluation period, the rational and life-like quality of the dolphins were confirmed by the biologists and ethologists, and also by the visitors. However, the emotions were not always clearly perceived. Additionally, the fact that most of the visitors tried to press several

sensors at the same time without waiting for explicit feedback did not help in understanding the IVE. For a visitor to understand the personality of the actors, she has to take some time to observe the dolphin interaction with its surroundings and other characters or, at least, be briefed about the dolphins' habits and behavior. In the context of a world exposition pavilion with no guided tour and a high visitor throughput, this was rarely the case. The average visitor spent an average of 2 minutes in the room. Thus, if the 10 people supervising the exposition rooms of the pavilion perfectly understood the dolphin behavior in the first two weeks, the same cannot be said about the average visitor who spent rarely more than a couple of minutes in the room.

To overcome this problem, a set of high priority rules overrode the dolphin behavior when sensor information was gathered. For instance, even if a dolphin had a high need for oxygen, if a visitor pressed the pleased sensor, it would stop emerging from the water and perform a typical pleased behavior (e.g., a loop). This was as far as we could go to convey immediate perceivable behavior without ruining the realism of the simulation and without compromising the project philosophy. The behavior interruption was not perceivable by the visitor since it is very hard for a new visitor to guess the design of a swimming dolphin. The same was obviously not the case for the supervisors.

Several alternatives were considered. A textual feedback of what was going on in the IVE could remove all the ambiguity of the visual feedback, but this was ruled out, mainly because of the multilinguistic nature of Expo'98. Another solution was to use posters to explain some typical behaviors, but it was found out that the average visitor would not bother reading them.

Considering the visitor perspective, the dialectic between artistic and real-world simulation was the major factor influencing and explaining the obtained results.

CONCLUSIONS

An architecture for rapid development of IVEs and the implementation of the mind module, controlling the agents emotional behavior, have been described. The model was tested on a specific application, S3A, a full real-time 3D IVE featuring two emotional dolphins, Tristão and Isolda, and was exhibited during the length of Expo'98 to more than a million visitors and developed under a tight schedule.

The three-part modular architecture implementing the theatrical metaphor (script–actors–stage) was adequate for the short development time of S3A. Its communication paradigm allowed heterogeneous teams from different areas of expertise to develop area-specific techniques to solve particular

problems efficiently, without worrying too much about the synchronization with the other teams. Indeed, the development (from scratch), took less than four months and ran at a minimally acceptable frame rate on a midrange graphical system.

The application also confirmed the adequacy of the conceptual mind abstraction for short development time of pathematic agents. Although the aspect of emotional believability was not undoubtedly perceived by all visitors, due to the characteristics of Expo'98, the architecture confirmed its suitability for implementing life-like believable characters with predefined personalities. The mind module implementation was based on the assumption that personality and emotions are essentially the same mechanism. Such an assumption allowed the characters' personality to be specified by a set of emotional components based on the OCC cognitive theory of emotions. This methodology proved to be an intuitive definition process and, by providing a common ground, allowed a heterogeneous team to create S3A. The four-phase character cycle (perceive-react-reason-act), inspired by Frijda's theory of emotions, was adequate to manage the emotional aspects of each character.

Indeed, S3A showed that the approach followed is a viable solution for short-time development of quality intelligent virtual environments.

NOTES

1. The diagram uses the unified modeling language notation (Fowler & Scott, 1997).
2. Improvisation directives are strings in S3A.
3. Team Technology, "Myers-Briggs: the Dynamic Model," http://www.teamtechnology.co.uk.
4. International Personality Item Pool, "Big-Five Factor Markers Derived from the IPIR," http://ipip.ori.org/ipip.
5. Nevertheless, the statistics for emotional behavior were very delicate to deal with, since what is really important is that the value of the behavior features at a determined point of time will provoke some very particular behaviors.

A EMOTIONAL REACTIONS (SIMPLIFIED)

A.1 Example 1- Attraction Reaction

"Tristão, our male dolphin, hates objects of human fabrication. When Tristão gets near a piece of the plane crash, he feels hate towards it."

	class:	attraction reaction
	valence:	negative appeal

active reaction	implementation:	`( [gen77] of RATTRACTION`
		`(agent TRISTAO)`
		`(object DWRECK)`
		`(appeal -9)`
		`(... other OCC variables ...))`

event		entity SOUTH-WRECK (class DWRECK) in the same area.
	type:	aspect of object emotion
	class:	hate (negative appeal)

launched emotion	implementation:	`( [gen78] of EHATE`
		`(agent TRISTAO)`
		`(potential 9) ; (= appeal)`
		`(intensity 4) ; (threshold = 5)`
		`(cause SOUTH-WRECK)`
		`(direction SOUTH-WRECK))`

A.2 Example 2- Attribution Reaction

"As Isolda, its female companion, approaches such hideous artifacts, he admires her courage."

	class:	attribution reaction
	valence:	positive praiseworthiness

active reaction	implementation:	`( [gen77] of RATTRIBUTION`
		`(agent TRISTAO)`
		`(subject DDOLPHIN)`
		`(id STANDARD-EXPLORE)`
		`; Standard of behaviour`
		`(objects DWRECK)`
		`(parameters ANY)`
		`(location ANY)`
		`(praiseworthiness +7)`
		`(... other OCC variables ...))`

event	match	`( [gen78] of EVENT`
		`(subject ISOLDA) ; of class DDOLPHIN`
		`(id EXPLORE)`
		`(objects SOUTH-WRECK) ; of class DWRECK`
		`(parameters veryClose)`
		`(location LSOUTHEAST-BOTTOM))`

	type:	self attribution emotion
	class:	admiration (positive praiseworthiness)

launched emotion	implementation:	`( [gen79] of EADMIRATION`
		`(agent TRISTAO)`
		`(potential 7) ; (= praiseworthiness)`
		`(intensity 3) ; (threshold = 4)`
		`(cause [gen78])`
		`(direction ISOLDA))`

A.3 Example 3- Event Reaction

"Isolda experiences great joy while exploring."

	class:	event reaction
	valence:	positive desirability

active reaction	implementation:	`( [gen77] of RWELL-BEING`
		`(agent ISOLDA)`
		`(subject DDOLPHIN)`
		`(id STANDARD-EXPLORE)`
		`; Standard of behaviour`
		`(objects DWRECK)`
		`(parameters ANY)`
		`(location ANY)`
		`(desirability +7)`
		`(... other OCC variables ...))`

event	match	`( [gen78] of EVENT`
		`(subject ISOLDA) ; of class DDOLPHIN`
		`(id EXPLORE)`
		`(objects SOUTH-WRECK) ; of class DWRECK`
		`(parameters veryClose)`
		`(location LSOUTHEAST-BOTTOM))`

	type:	prospect irrelevant emotion
	class:	joy (positive desirability)

launched emotion	implementation:	`( [gen79] of EJOY`
		`(agent ISOLDA)`
		`(potential 7) ; (= desirability)`
		`(intensity 4) ; (threshold = 3)`
		`(cause [gen78])`
		`(direction NIL))`

A.4 Example 4- Prospect Reaction I

"When Tristão's stomach takes control over his body, he starts a new goal: GEat. He feels hope as likelihood is high– he rarely misses his prey."

	class:	prospect reaction
	valence:	positive (high likelihood)

active reaction	implementation:	`( [gen77] of RPROSPECT`
		`(agent TRISTAO)`
		`(goal GEAT)`
		`(... OCC variables ...))`
		where
		`( [GEAT] of GACTIVEPURSUIT-STANDARD`
		`(agents DDOLPHIN)`
		`(importance-success 4)`
		`(importance-failure 4)`
		`(likelihood GEat-efn)`
		`; assume 75% this cycle`
		`(... other OCC variables ...))`

event		change of likelihood (activation) of GEat from 0% to 75%
	type:	prospect relevant emotion
	class:	hope

launched emotion	implementation:	`( [gen79] of EHOPE`
		`(agent TRISTAO)`
		`(potential 7.5) ; (= likelihood*10)`
		`(intensity 4.5) ; (threshold = 3)`
		`(cause GEat)`
		`(direction NIL))`

A.5 Example 5- Prospect Reaction II

"When he finally gets his prey, he feels satisfaction."

	class:	prospect reaction
	valence:	positive (existing hope)

active re-action	implementation:	`( [gen77] of RPROSPECT`
		`(agent TRISTAO)`
		`(goal GEAT)`
		`(... OCC variables ...))`
		where
		`( [GEAT] of GACTIVEPURSUIT-STANDARD`
		`(agents DDOLPHIN)`
		`(importance-success 4)`
		`(importance-failure 4)`
		`(likelihood GEat-efn)`
		`(... other OCC variables ...))`

event		success of goal GEat.
	type:	prospect relevant emotion
	class:	satisfaction

launched emotion	implementation:	`( [gen79] of ESATISFACTION`
		`(agent TRISTAO)`
		`(potential 4) ; (= importance of success)`
		`(intensity 1) ; (threshold = 3)`
		`(cause GEat)`
		`(direction NIL))`

REFERENCES

Ball, G., D. Ling, D. Kurlander, J. Miller, D. Pugh, T. Skelly, A. Stankosky, D. Thiel, M. van Dantzich, and T. Wax. 1997. Lifelike computer characters: The persona project at Microsoft Research. *Software Agents*, ed. J. Bradshaw 191–222. Menlo Park, California: AAAI/MIT Press.

Bates, J. 1994. The role of emotion in believable agents. *Communications of the ACM* 37(7): 122–125.

Badler, N., B. Reich, and B. Webber, 1997. Towards personalities for animated agents with reactive and planning behaviours. *Creating Personality for Synthetic Actors*, eds. R. Trappl and P. Petta, 43–57. New York: Springer–Verlag.

Blumberg, B., and T. Galyean. 1995. Multi–level direction of autonomous creatures for real–time virtual environments. In *SIGGRAPH '95 Computer Graphics*, *Proceedings of the 22nd Annual ACM Conference*, 47–54. New York: ACM Press.

Brooks, R. 1986. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* 2(1): 14–23.

Cañamero, D. 1997. Modelling motivation and emotions as a basis for intelligent behavior. In *Proceedings of the First International Conference on Autonomous Agents*, eds. W. L. Johnson and B. Hayes–Roth, 148–155. New York: ACM Press.

Damásio, A. 1994. *Descartes' Error*. Publicações Europa–América, Mem Martins, Portugal.

Elliot, C., J. Lester, and J. Rickel. 1997. Integrating affective computing into animated tutoring agents. *In Working Notes of the International Joint Conference on Artificial Intelligence '97 Workshop on Animated Interface Agents*, 113–121.

Fowler, M., and K. Scott. 1997. *UML Distilled: Applying the Standard Object Modelling Language*. Addison–Wesley Longman.

Frijda, N. 1987. *Emotions*. London: Cambridge University Press.

Geib, C., L. Levison, and M. Moore. 1994. *SodaJack: An Architecture for Agents that Search for and Manipulate Objects*. Technical Report MS–CIS–94–16, University of Pennsylvania, Philadelphia, PA.

Grand, S., D. Cliff, and A. Malhotra. 1997. Creatures: Artificial autonomous software agents for home entertainment. In *Proceedings of the First International Conference on Autonomous Agents*, eds. W. L. Johnson, and B. Hayes–Roth. New York: ACM Press.

Hayes–Roth, B., and R. van Gent. 1996. *Story–Making with Improvisational Puppets and Actors*. Technical Report KSL–96–05, Stanford University, Stanford, CA.

Hayes–Roth, B., and D. Rousseau. 1996. *Personality in Synthetic Actors*. Technical Report KSL–96–28, Stanford University, Stanford, CA.

Hayes–Roth, B., R. van Gent, and D. Huber. 1996. Acting in character. *Creating Personalities for Synthetic Actors*, eds. R. Trappl and P. Petta, 92–112. New York: Springer–Verlag.

Johnson, L., J. Rickel, R. Stiles, and A. Munro. 1998. Integrading pedagogical agents into virtual environments. *Presence* 7(6): 523–546.

Lester, J., and B. Stone. 1997. Increasing believability in animated pedagogical agents. In *Proceedings of the First International Conference on Autonomous Agents*, eds. W. L. Johnson and B. Hayes–Roth, 16–21. New York: ACM Press.

Lester, J., S. Converse, S. Kahler, T. Barlow, B. Stone, and R. Bhoga. 1997a. The personal effect: Affective impact of animated pedagogical agents. In *Proceedings of Computer Human Interaction 97 Conference*, ed. S. Pemberton, 359–366. Reading, MA: ACM/Addison–Wesley.

Lester, J., J. Voerman, S. Towns, and C. Callaway. 1997b. Cosmo: A life–like animated pedagogical agent with deictic believability. In *Working Notes of the International Joint Conference on Artificial Intelligence '97 Workshop on Animated Interface Agents*, 61–69.

Loyall, A., and J. Bates. 1993. Real–time control of animated broad agents. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society*.

Maes, P., T. Darrell, B. Blumberg, and A. Pentland. 1997. The ALIVE system: Wireless, full–body interaction with autonomous agents. *Multimedia System*, 5(2): 105–112.

Martinho, C. 1999. *Emotions in Motion*, M.Sc. Thesis, Technical University of Lisbon, Portugal, IST.

Moffat, R. 1997. Personality parameters and programs. *Creating Personalities for Synthetic Actors*, eds. R. Trappl and P. Petta, 120–165. New York: Springer–Verlag.

Ortony, A., G. Clore, and A. Collins. 1988. *The Cognitive Structure of Emotions*. London: Cambridge University Press.

Paiva, A., and I. Machado. 1998. Vincent, an autonomous pedagogical agent for on-the-job training. *Intelligent Tutoring Systems*, eds. B. Goettl, H. Halff, C. Redfield, and V. Shule. New York: Springer–Verlag.

Perlin, K. 1996. Improv: A system for scripting interactive actors in virtual worlds. In: *SIGGRAPH '96 Computer Graphics*, *Proceedings of the 23rd Annual ACM Conference*. New York: ACM Press.

Reilly, W. 1996. *Believable Social and Emotional Agents*, Ph.D Thesis, Carnegie Mellon University, Pittsburgh, PA.

Rousseau, D., and B. Moulin. 1996. *Personality in Computer Characters*, Technical Report KSL 96-13, Stanford University, Stanford, CA.

Rousseau, D., and B. Hayes-Roth. 1996. *Personality in Synthetic Agents*, Technical Report KSL-96-21, Stanford University, Stanford, CA.

Sepúlveda, T. 1996. *Synthetic Actors in Virtual Environments*, M.Sc. Thesis, Technical University of Lisbon, Portugal.

Terzopoulos, D., X. Tu, and R. Grzeszczuk. 1994. Artificial fishes: Autonomous locomotion, perception, behaviour, and learning in a simulated physical world. *Artificial Life* 1(4):327–351.

Thalmann, D., H. Noser, and Z. Huang. 1997. Autonomous virtual actors based on virtual sensors. *Creating Personality for Synthetic Actors*, eds. R. Trappl and P. Petta, 25–42. New York: Springer–Verlag.