The Nature of Characters in Interactive Worlds and The Oz Project

Joseph Bates

October 1992

CMU-CS-92-200

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

To appear in Virtual Realities: Anthology of Industry and Culture, Carl Eugene Loeffler, Editor, 1993.

Abstract

Traditional storytelling media, such as the novel, cinema, and television, draw much of their emotional power from characters and story. The greatest impact of interactive media will arise only when they too exhibit these characteristics. However, rich interactive characters and rich interactive story are unfamiliar concepts. The Oz Project is trying to bring together writers, artists, and artificial intelligence researchers to produce fundamental technology that can support this new form of art and entertainment. We discuss these ideas, argue for the essential role AI has to play in interactive media, and present a non-technical introduction to the construction of interactive characters.

Keywords: artificial intelligence, autonomous agents, interactive characters, interactive media, virtual reality, simulated worlds, Oz

1 The Demands of Interactivity

Virtual reality is the most recent form of interactive world. Earlier examples include the video games that were perhaps first invented in the 1960's and the text adventures that were the most popular home software in the early 1980's. However, virtual reality has touched the hearts of many people in ways that earlier simulations did not.

A result of this great popularity is widespread speculation on the future of virtual worlds. A frequent notion in such speculation is that we will soon develop "interactive cinema", which will allow people to enter into immersive environments having as much emotional impact as, say, *Citizen Kane* or *Terminator 2*. After all, we can already build simulated physical spaces and simulated physical objects, so we must be about half of the way toward including simulated characters and an exciting story.

This speculative notion, while most attractive, is flawed. Consider the simulated characters. Writing about and filming characters is not fundamentally more difficult than writing about and filming scenery. So for non-interactive media, you are indeed half-way there when you can manage scenery. In contrast, developing richly interactive characters means constructing intelligent, emotional, behaving creatures. These creatures must seem to live in the simulated world and must respond to the user's rich variety of human behavior in believable ways. This kind of creature is a primary goal of basic research in artificial intelligence, and it is not merely a next obvious step for designers of virtual reality systems.

Similarly, interactive story is not an obvious extension of traditional linear stories. In order to feel the reality of a simulated world, users must be free to behave in any ways that fit within the theme of the world. For instance, a flight simulator need not support stopping for lunch at a McDonald's, but it must permit the freedom to use the airplane controls in any reasonable way. However, a story is by nature an imposition of structure on the user, because whatever the user chooses to do must inevitably lead to the dramatic arc of some story-like experience. Thus, an interactive story system must provide a computational solution to the apparent clash between free will and destiny. This is unlike the requirements of story in traditional media, and is more than a minor extension to existing virtual reality technology.

If one views a VR system as producing surface level phenomena via the hardware interface and associated software, then the organization and content of software well behind the interface constitute a "deep structure" for the virtual world. This is the arrangement of code and data that produces the central meanings of the interactive experience. For the reasons suggested above, we believe that for VR to join the novel, cinema, and television as a broadly successful artistic medium, the technology must provide a sufficiently rich deep structure. In particular, it must provide computational theories for interactive characters and story.

The Oz project at Carnegie Mellon is our attempt to work toward these goals.

The project includes artists, writers, and researchers in artificial intelligence and computing. We believe that such a collaboration is necessary to capture the artistic knowledge, sensibilities, and capabilities needed to provide a deep structure for dramatic worlds.

Current Oz research falls into three areas:

- the construction of broadly capable, though perhaps shallow, autonomous agents that integrate elements of perception, cognition, emotion, action, and language,
- the construction of a computational theory of interactive drama, to gently shape the user's overall experience, and
- the development of computational methods for varying the presentation style of the experience, thus providing the interactive analog of film technique and writing style.

Our work is described in detail in a number of reports [1, 2, 3, 4, 5, 6]. The rest of this note summarizes our research on the construction of agents.

2 Interactive Characters

One of the keys to an effective virtual world is for the user to be able to "suspend disbelief". That is, the user must be able to imagine that the world portrayed is real, without being jarred out of that belief by the world's behavior. In existing works of interactive fiction and other simulated worlds, the unnatural behavior of simulated agents is perhaps the primary impediment to fully suspending disbelief.

Traditional research on agents in artificial intelligence demands that constructed creatures be highly competent. Our central requirement, that users be able to suspend disbelief, is different and unusual. Instead of demanding that agents be especially active and smart, we require only that they not be clearly stupid or unreal. An agent that keeps quiet may appear wise, while an agent that oversteps its abilities will probably destroy the suspension of disbelief. Thus, in Oz we try to take advantage of the "Eliza effect" [7], in which people see subtlety, understanding, and emotion in an agent as long as the agent does not actively destroy the illusion.

In order to foster this illusion of reality, we believe our agents must have broad, though perhaps shallow, capabilities. To this end, we have developed an architecture for mind, called Tok, that exhibits some signs of internal goals, reactivity, emotion, natural language ability, and knowledge of agents (self and other) as well as of the simulated physical micro-world.

Rather than describing Tok in detail (the earlier mentioned technical papers provide such information), we will convey the flavor of our work by presenting a particular

Emotions	Behaviors		Features
hope†	wanting to be pet *or brushed*	purring	curious
fear	cleaning self	arch back	content
happy	wanting to go out/in	hiss	aggressive
sad	eating	swat	ignoring
pride	wanting to eat	bite	friendly
shame	getting object (using human	escape/run away	*proud*
admiration	or other tool)	have fun	energetic
reproach	searching for something	*pouncing on creatures*	
gratification	*carrying mouse*	chasing ball/*creatures*	
remorse	playing with ball	rubbing against	
gratitude	playing with mouse	licking	
anger	*crazy hour*	watching/staring at	
love	hiding (anger/fear)	sitting on a sunny ledge	
hate	pushing things around		

†*italicized items were not included in final implementation*

Table 1: Original Lyotard Task

agent built in Tok. That agent is a simulated house cat named "Lyotard". Our goal in developing Lyotard was to build a creature that could believably pass for a cat in an Oz micro-world.

Table 1 lists the emotions and behaviors from our original informal design document for Lyotard. The emotions are those naturally available in the current version of Tok, though in the end we did not use all of them. The behaviors were developed over several hours of brainstorming by several cat owners in our group. The behavioral features are used to modify the details of Lyotard's behaviors, according to Lyotard's mood.

3 The Behavior of Lyotard

To our knowledge, whether an agent's behavior produces a successful suspension of disbelief can be determined only empirically. The agent must be embedded in a world, and a variety of users must report their subjective experience with the agent. For us this evaluation is an on-going effort, which we will attempt to report in the literature [5] and to convey by demonstration.

Nonetheless, we hope to provide the reader of this non-interactive text with some sense of Lyotard's behavior and of the internal mental causes for the behavior. Thus, we present in Table 2 a small excerpt of a session with Lyotard. In this session a human user interacted with Lyotard in a simulated six room house. Because we are interested in the actions of the agents, the figure presents the actions of each agent from an omniscient perspective. The normal output from the Oz system to the human user, such as English descriptions of what the human perceives, prompts for the human's action, etc., have been omitted.

```
Lyotard:  (*GO-TO "the bedroom").
          (*GO-TO "the sunroom").
          (*GO-TO "the spare room").
          (*JUMP-ON "the chair").
          (*SIT-DOWN).
          (*LICK "Lyotard").
          (*LICK "Lyotard").

Player:   (*GO-TO "the spare room").
Lyotard:  (*JUMP-OFF "the chair").
          (*RUN-TO "the sunroom").

Player:   (*GO-TO "the sunroom").
Lyotard:  (*LOOKAROUND NERVOUSLY).
Player:   (*PET "Lyotard").
Lyotard:  (*BITE "Player").
          (*RUN-TO "the diningroom").

Player:   (*GO-TO "the spare room").
Lyotard:  (*LOOKAROUND NERVOUSLY).
          (*GO-TO "the sunroom").
          (*POUNCE-ON "the superball").
          (*LOOKAT "the superball").
          (*NUDGE "the superball").
          (*POUNCE-ON "the superball").
          (*POUNCE-ON "the superball").

Lyotard:  (*GO-TO "the diningroom").

Lyotard:  (*GO-TO "the kitchen").
          (*MEOW).
Player:   (*GO-TO "the sunroom").
Lyotard:  (*MEOW).
Player:   (*GO-TO "the diningroom").
Lyotard:  (*WAIT).
Player:   (*TAKE "the glass jar").

Lyotard:  (*GO-TO "the diningroom").
Player:   (*GO-TO "the kitchen").
Lyotard:  (*JUMP-ON "the table").

Lyotard:  (*JUMP-OFF "the table").
          (*GO-TO "the kitchen").
          (*MEOW).
Player:   (*POUR "the glass jar" IN
                 "the kitty bowl").

Lyotard:  (*EAT "the sardine").
          (*EAT "the sardine").
          (*EAT "the sardine").
          (*EAT "the sardine").
          (*EAT "the sardine").

Player:   (*PET "Lyotard").
Lyotard:  (*CLOSE-EYES LAZILY).
Player:   (*TAKE "Lyotard").
Lyotard:  (*CLOSE-EYES LAZILY).
```

Table 2: Section of an interaction with Lyotard

Just prior to the beginning of this excerpt, Lyotard had successfully finished exploring part of the house. This success made Lyotard mildly happy. The happy emotion led to the content feature being set, which gave rise to a behavior to find a comfortable place to sit. In finding a comfortable place to sit, Lyotard remembers places that he believes to be comfortable and chooses one of them, a particular chair in the spare room. He then goes there, jumps on the chair, sits down, and starts cleaning himself for a while.

At this point, the human user, whom Lyotard dislikes, walks into the room. Tok models like and dislike as relatively long-lasting attitudes toward external entities. The initial mild dislike of the user gives rise to an emotion of mild hate toward the user. Further, Tok notices that one of Lyotard's goals, not-being-hurt, is threatened by the disliked user's proximity. This prospect of a goal failure generates fear in Lyotard. The fear and hate combine to generate a strong aggressive feature and to diminish the previous content feature. The fear emotion and the proximity of its cause give rise to an avoid-harm goal, while the aggressive feature gives rise to a goal to threaten the user. In this case the avoid-harm goal wins out, creating a subsidiary escape/run-away behavior that leads Lyotard to jump off the chair and run out of the room.

When the user follows Lyotard into the sunroom and tries to pet him, Lyotard sees the action and notices that the actor trying to touch him is one toward whom he feels mild hate. This combination generates another goal: respond-negatively-to-contact. Lyotard responds to this rather than to his escape/run-away goal or any of his other goals because we declared it as having a high priority when we created Lyotard. Further refinement of this goal through a series of choices leads to Lyotard biting the player.

As the player leaves Lyotard alone, the emotions engendered by the player start to decay, and Lyotard again pursues his amusement goal. This time he is no longer content, which is one of several changes to his emotional state, so a slightly different set of amusement choices are available. He chooses to play with one of his toys, and so goes to find his superball.

As the simulation has progressed, Lyotard's body has been getting more hungry. At this point his hunger crosses a threshold so that his mind notices it as a feeling of hunger. This triggers a feeding goal causing him to go to his bowl, but it is empty so he complains by meowing. After a while, he gives up on this technique for getting food, so he tries another technique; he goes looking for food himself. He remembers places where he has seen food that was reachable, and goes to one of them, passing by the user in the process. At this point he again feels fear and aggression, but he ignores these feelings because dealing with the hunger is more important to him. As he reaches the location he expected to find the food, he notices that it is gone (taken by the user when Lyotard couldn't see him), so Lyotard again considers other techniques to get food. He could try to find a human and suggest he be fed, but instead he chooses to try his bowl again. This time the human feeds him, and Lyotard eats. As he eats he feels happy because his emotionally important goal of eating is succeeding, and he also feels gratitude toward the user, because he believes the user helped to satisfy this goal. This gratitude in turn gradually influences Lyotard's attitude toward the user from dislike to neutral.

Now when the user pets Lyotard, Lyotard responds favorably to the action by closing his eyes lazily. Lyotard wants to be pet because he no longer dislikes or fears the user. Thus, being pet causes a goal success which causes happiness, and because the goal success was attributed to the user, increases gratitude toward the user. The result is that Lyotard now strongly likes the player.

The trace we have shown was produced by the interactive fiction version of Oz, which is written in Common Lisp. Of the 50,000 lines of code that comprise Oz, the Tok architecture is roughly 7500 lines. Lyotard is an additional 2000 lines of code.

Running on a Hewlett Packard 720 workstation (55 MIPS), each Tok agent takes roughly two seconds for processing between acts.

4 Conclusion

The VR community has generally focused on VR as a human-interface technique, with some attention also given to modeling physical space and objects. While these topics are important, we see exclusive attention to them as something like studying celluloid instead of cinema, paper instead of literature, or cathode ray tubes instead of television. To reach our dream of "interactive cinema", we must also look at the underlying content of the worlds we want to model. This means studying interactive characters, story, and presentation style, and that in turn means studying artificial intelligence.

The Oz project is an attempt to explore these areas, in which we blend AI technology with ideas and insight from traditional arts and media. Besides the Lyotard world, discussed above, we also have built a real-time, interactive, animated world containing several autonomous emotional/cognitive creatures. These "Woggles" were first shown at the AAAI-92 Artificial Intelligence Based Arts Exhibition, in San Jose, California. Other Oz work is in progress on interactive story and presentation [4, 5].

The long term goal for Oz is to develop a popular and widespread new form of art and entertainment. Oz simulations, which today require large engineering workstations, will run on personal computers (or television sets) of the middle and late 1990's. Such AI-based interactive entertainment software may be a key element driving the merger of computing and consumer electronics beyond this decade.

5 Acknowledgments

We gratefully acknowledge the support of Fujitsu Laboratories, Ltd. for our research.

References

[1] Joseph Bates. Virtual reality, art, and entertainment. *PRESENCE: Teleoperators and Virtual Environments*, 1(1):133–138, 1992.

[2] Joseph Bates, A. Bryan Loyall, and W. Scott Reilly. Broad agents. In *Proceedings of AAAI Spring Symposium on Integrated Intelligent Architectures*, Stanford, CA, March 1991. Available in *SIGART Bulletin*, Volume 2, Number 4, August 1991, pp. 38-40.

[3] Joseph Bates, A. Bryan Loyall, and W. Scott Reilly. Integrating reactivity, goals, and emotion in a broad agent. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Bloomington, IN, July 1992.

[4] Mark Kantrowitz and Joseph Bates. Integrated natural language generation systems. In R. Dale, E. Hovy, D. Rosner, and O. Stock, editors, *Aspects of Automated Natural Language Generation*, volume 587 of *Lecture Notes in Artificial Intelligence*, pages 13–28. Springer-Verlag, 1992. (This is the Proceedings of the Sixth International Workshop on Natural Language Generation, Trento, Italy, April 1992.).

[5] Margaret Kelso, Peter Weyhrauch, and Joseph Bates. Dramatic presence. Technical Report CMU-CS-92-195, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, October 1992. Submitted to PRESENCE: Teleoperators and Virtual Environments, MIT Press, Cambridge, MA.

[6] W. Scott Reilly and Joseph Bates. Building emotional agents. Technical Report CMU-CS-92-143, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, May 1992.

[7] Joseph Weizenbaum. Eliza. *Communications of the ACM*, 9:36–45, 1966.