Multiple Conceptions Of Character-Based Interactive Installations

Bruce Blumberg Bill Tomlinson Marc Downie
Synthetic Characters Group / The Media Lab, MIT
Room E15-309 / 20 Ames Street
Cambridge, MA 02139 USA
+1 617 253 5109
{bruce | badger | marcd}@media.mit.edu

ABSTRACT
This paper presents ways of approaching the design of successful character-based interactive installations. We rationalize our arguments within the context of both Disney's "illusion of life" and Daniel Dennett's "intentional stance". We present six perspectives from which intentional characters can be viewed: as interactors on a variety of time scales; as reciprocal interactors with each other; as entities exhibiting a dynamic expressive range; as creatures with life cycles; as a combination of allusions to existing media; and as a collection of well-balanced components. By conceptualizing characters in these ways, creators can generate installations that enable participants to read the desires, beliefs, and actions of the characters. This approach forms the basis of a successful character-based interactive installation.

Keywords
Animated characters, interactive graphics, interface design, autonomous agents.

1. Introduction

Character animation is fundamentally the art of revealing a character's inner thoughts (its beliefs and desires) through motion, sound, form, color and staging. As Disney animators Frank Thomas and Ollie Johnson [15] put it their book *The Illusion of Life*, "It is the change in shape that shows what the character is thinking. It is the thinking that gives the illusion of life. It is life which gives meaning to the expression." When we look at a great animated character we can tell what that character is thinking and feeling, and while we may not know exactly what a character is about to do, we have a pretty good idea based on our perception of its desires and beliefs. Even when our guess is wrong, the resulting behavior is almost always explicable.

While the "illusion of life" makes it clear what one must do if one wants to bring a character to life, it does not address the question of why these techniques work. A concise explanation can be found in the work of philosopher Daniel Dennett. Dennett argues that we take the "intentional stance" [4] with respect to predicting and explaining the actions of animate things in our world, including people, animals, and animated characters. The intentional stance is pretty simple. First, you start with a character's desires. Then you figure out what its beliefs about the world should be, given its situation in the world. Finally, you assume that it will act in a "character-specific" way so as to satisfy those desires given its beliefs, and given the state of the world. Seen in this way, we use the intentional stance to predict a character's actions based on our knowledge of its presumed desires and beliefs. Most importantly, we use *our* belief that the character's motion and the quality of that motion is a direct consequence of its underlying desires and beliefs is used to infer a character's desires and beliefs. Indeed, this is the point that Thomas and Johnson are making in their quote. When one looks at the "illusion of life" in the context of Dennett's work, one sees that the techniques put forth are essentially a recipe for making it easy for the viewer to take the intentional stance toward a character.

The challenge for a control system for an autonomous animated character can be put in equally simple terms: make it easy for the participants to take the intentional stance with respect to that character. Thus, the control system must ensure that the character's motion and quality of motion are a direct and clear consequence of the character's desires and beliefs, and conversely that the its desires and beliefs can be easily inferred from its resulting motion and quality of motion. Of course, the character's desires and beliefs, and how those desires and beliefs change over time, must make sense given who the character is and given the character's interactions, past and present, with its world.

The challenge for an installation – everything else that surrounds the characters – is to let people read the

intentionality of the characters. Indeed, much of the wisdom in the illusion of life centers on how to use staging, camera, lighting, color, sound and music to enhance the observer's understanding and connection with the desires, beliefs and actions of the characters. But how does one do this in an interactive installation in which the characters have minds of their own, the observer is in fact a participant, and the interaction must be compelling and believable over multiple time scales?

In this paper we present a variety of ways to conceive of character-creation that we believe help ensure that it is both easy and interesting for the participant to take the intentional stance toward the characters. The first of these ways is to analyze the interactions of a character on a range of time scales, ensuring that the character behaves appropriately at each level. The next is to consider the interactions that each character has with every other element in the installation (including other characters, human participants, cinematography systems, dynamic musical scores, etc.) The third is to consider the behavioral dynamic range of the character; characters should be capable of a wide range of expressive behavior, even if they tend to stick to one area of that range most of the time. A fourth way to view a character is as a life cycle. Just as animals undergo growth and development over the course of their life, it is important to consider how an autonomous character will change over the course of its life. A fifth way to consider a character is as a unified individual, in whom all the various components work well together and integrate closely. A final perspective from which to think about a character is as a combination of allusions to existing media with which the participant is already familiar. The unifying theme throughout is the question: how does the installation as a whole help support taking the intentional stance toward the characters?

Tweedie [18] proposed a means of characterizing interactive techniques. Her work sought to understand the relationships between inputs and outputs in visualizations of data. However, she did not address installations that feature characters. Installations that feature characters are a significant venue for the exploration of human computer interaction. Rickenberg [12] studies the usefulness of animated characters and how they affect task oriented human-computer interactions. Installations with characters for entertainment call for a different approach. Pandzic et al [8] have done very impressive work with distributed interactive environments.

Each way of thinking about characters does not conflict with the others; rather, it provides an additional perspective from which to analyze characters, and thereby installations. As Seymour Papert suggests, it can be valuable to have multiple "ways of knowing" a topic. [9]

2. Installations

Bruce Blumberg's Synthetic Characters Group at the MIT Media Lab has created several large-scale interactive installations [1,2,17] in recent years and an assortment of smaller pieces that we have shown to visitors to our lab. Since many of the sections of this work will be illustrated with examples from our installations, we begin by giving a short summary of each of the pieces that we have done over the last several years.

Swamped! [1] featured a virtual chicken running around a barnyard scenario on a large projection screen. Participants could control the chicken by means of a plush toy (a bright yellow fleece chicken). A fully autonomous raccoon marauded around the barn yard in search of the chicken's eggs. The *Swamped!* installation featured an action selection mechanism, a motor system and a novel interface for interacting with autonomous and semi-autonomous characters. [6]

(void): A Cast of Characters* [2] showed three humanoid characters sitting at an all-night diner. The interface, two dinner rolls with forks stuck in them (in the spirit of Charlie Chaplin's film "The Gold Rush"), allowed a participant to make the characters get up and dance. The characters were different from puppets, though, having emotional responses to the interaction that they were undergoing and changing the entire style of their animation and interaction to reflect their emotional state. In addition, the characters in *(void*)* could learn the ways in which people interacted with them, and would continue to act in those ways after the participant stopped interacting. [19,20]

sand:stone [17] was an interactive art installation based on the poem "Ozymandias" by Percy Bysshe Shelley. In this piece, people could moves stones around on a surface of sand. The positions and relationships of the stones caused changes in a projected display of an animated statue of a great king. The installation explored notions of decay and the passage of time.

Finally, two of our latest projects feature an adaptive autonomous animated terrier named *Duncan. sheep|dog: Trial by Eire* showed at the opening of the

MediaLabEurope in Dublin, Ireland in July 2000, and allowed participant to play the role of a shepherd in a virtual sheep herding competition. By means of his trusty terrier the participant was able to coax a flock of ornery autonomous sheep around a field and into a pen. In *Clicker By Eire*, the participant, using dog training techniques borrowed from the real world, may train *Duncan* to perform a variety of tricks in response to verbal cues. Technically, *Duncan* is our platform for focusing further on learning, action selection and motor control in autonomous characters.

In the coming sections, we discuss various topics that have come up during the creation of these installations from the point of view of a several person team putting together a fairly large-scale computer graphical interactive installation with some sort of tangible interface. Many of the topics, though, are relevant to other pursuits that deal with the juncture between people and technology.

3. Interaction Time Scales

In order to make sure that a character is interesting in more than just a shallow way, we analyze its behavior in terms of *interaction time scales*. We look at four time scales of interactivity and illustrate them with examples from characters in our installations:

t < 2 seconds: In the very short term, a character should be *responsive*. It should react quickly to any occurrence that a participant would expect it to be able to sense (a prime example being any kind of input from the participant). The *(void*)* characters looked down at their feet as soon as a participant began moving the interface, which let people know that the characters were not only aware of them but also were expecting their feet to be moved as a result. If the characters are not responsive, the installation runs the risk of being labeled *broken*.

2 seconds < t < 30 seconds: On a slightly longer time scale, a character should be *understandable*. Once responsiveness is verified, people want to be able to understand *what* the character is doing. *(void*)* characters were understandable: when a participant wiggled the rolls in a certain way, the character on screen danced in a similar way. If characters fail on this time scale, the installation is *confusing*.

30 seconds < t < 15 minutes: After the initial novelty of an installation has worn off, characters need to be *interesting* in order to hold the attention of participants. *(void*)* was interesting because as a participant interacted with the character, she had an effect on the emotional state of that character, which was displayed via multiple expressive channels (see section on Dynamic Range, below). Also, that character often continued to dance for a few moments, in whatever style the character had learned from the participant. If the characters are not interesting, the installation is *boring*.

15 minutes < t: Finally, in order for a piece to be engrossing at a deeper level, the characters and the interaction must be *subtle*. Perhaps the most important way to achieve this is by using longer-term learning and adaptation. By taking dogs as the model for our most recent work (e.g. *sheep|dog*) we seek to address the creation of long term relationships, characters who recognize participants and are recognized by them. These subtle but lasting interactions will hopefully take installations beyond the 15-minute mark. If the characters are not subtle, the installation will eventually turn out to be *shallow*.

The *(void*)* characters were fleshed out to satisfy each of these four time scales. When a person began interacting with one of them, the character immediately looked at its feet to show that it was aware that something was happening. Soon thereafter, it began to move in time with the participant's moving of the buns, reflecting the actions that the participant was taking. After some time at this, the participant noticed the emotional change that the character was undergoing as a result of the interaction. Finally, over a longer time period, participants could experiment with the learning that the characters were able to do. In order for a virtual character to be seen as intelligent, it needs to behave in a rational fashion in each of the four time scales. Indeed, a good measure of the design of the underlying behavior system is the extent to which these interaction time scales emerge naturally.

4. Reciprocal Interactions

These interaction time scales must be properly grounded in the capabilities of the characters themselves. If the character's action selection mechanism cannot react to sudden loud noises, loud noises should not be part of the installation. The idea of matching the abilities of the characters with the complexities of their virtual worlds will appear throughout this paper.

A correlate of this rule is that no elements should be introduced to the installation that does not support the central interaction. For example, in *Swamped!*, one of our modelers made a beautiful tree house that looked great in the setting of this swamp. However, every

time we showed the developing installation to people, they immediately wanted to have their character run over and climb into the tree house. Since the installation was not meant to focus on the tree house (which was intended to be merely "eye candy"), we ultimately removed it. The only purpose it served was to defocus the characters' beliefs, intentions and desires, and to highlight their shortcomings instead.

A similar principle applies to inter-character relationships. Every character in an installation should be able to interact with everything that it encounters in its world – characters, objects, and human participants. Every character that cannot recognize the presence of another character endangers the illusion of presence.

The characters in *sheep|dog* provide a good example of this ability. The sheep generally try to move away from the dog and the shepherd. If the dog or shepherd gets too close, though, the sheep may charge at them. If the sheep charge the dog, the dog often gets scared and runs away. If they charge the shepherd, he jumps back. The dog responds to the shepherd's commands. The shepherd looks at the dog and gestures to him. Each of the characters has an active interaction with the others.

The participant, too, should be seen as one of the creatures participating in these interactive "dyads" [7]. In order to enhance this perspective, the salient elements of human social relationships may be brought to bear. For example, one of the greatest moments in the development of *Swamped!* was the first time we enabled the raccoon to look at the camera. This gave participants the distinct impression that it was making eye contact with them. (It is interesting to note that, even though participants interacted with *Swamped!* by means of a physical representation of the virtual chicken, they felt that the raccoon was looking at them when he looked at the *camera*, rather than at the virtual chicken.)

The relationship between the cinematography system (a "camera creature") and the characters in *Swamped!* is another example of a reciprocal interaction: the characters could look at the camera, and the camera was able to frame shots around the characters in an intelligent manner. We used this to even greater advantage in *(void*)*, in which all three characters would occasionally glance at the camera at significant moments. [16]. Looking at the participant creates a powerful feeling that the characters are aware of the participant, closing the interaction loop between participant and character.

(void)*'s characters were accompanied by a "music creature". Its goal: to create a film score in real time for a medium without a script. In a film, the score's entire purpose is to support what happens on screen – to help stage the emotional changes that take place within the characters and to give insight into their actions. A score, dynamic or not, that fails to do just that detracts from the piece as a whole. Hence the music creature had to be highly aware of both the characters on screen and the camera.

5. Dynamic Range

In film making, the best way to make a scene seem dark is to place a single small light source somewhere in an otherwise dark frame. This light defines the high end of the dynamic range of illumination. Just like a scene from a movie, a character should have a dynamic range of experience. Rather than leaving this range to chance, it is useful to intentionally define this dynamic range, and thereby control the experience that participants have in interacting with that character. For example, in (void*), we paid special attention to how much a participant had to do in order to provoke a character enough for that character to storm out of the diner, or to crack a big smile. A character that is grumpy all the time seems "flat". However, if that character can be coaxed into smiling (almost "against his will"), the character will seem much more convincing, and the installation featuring that character will be more interesting.

We have used two main emotional models: one works with six canonical emotions (Sadness, Happiness, Fear, Disgust, Surprise and Anger) [5]; the other, with three orthogonal axes (Stance, Valence and Arousal) [14]. Regardless of the internal representation, there are many ways to show how a character is feeling. The most obvious of these is facial animation. A smile is the clearest way to show that a character is happy. Body posture is very important, as is the way a character moves (fast, jerky motion vs. slow, smooth motion). What actions a character takes, in addition to how he takes them, can differentiate among emotions (a fearful character might flee if another character surprises him, whereas a happy character might just startle). Vocalizations are another powerful way to communicate emotional state.

It is worth noting that there is little point to having a sophisticated behavior system or emotional model if the underlying motor system is incapable of generating the quality of motion required to convey the characters' behaviors or emotions. In our work, for example, we use a verb-adverb model inspired by the

work of Rose [13] for specifying and controlling motion. In addition, we generally use key-framed animation as opposed to procedural animation as our source material. Through real time multi-target motion interpolation our motor systems are capable of generating a wide dynamic range of motion. A similar approach is taken by Perlin et al [10].

However, not all expressive channels lie within the functionality of the character itself. There are several extra-character mechanisms that we use to communicate that character's emotion to participants. The reactions of other characters can help giving the sense that the characters know each other and can perceive how each other is feeling, perhaps better than we can. Elements like this fill the installation with a greater feeling of character depth, and work well with the "Reciprocal Interactions" analytical technique described above. Dynamic cinematography and interactive lighting design [16] are useful for crossing the boundary between computers and people and conveying the characters' emotions to participants. For example, when the raccoon in *Swamped!* was very mad at the chicken, he would get reddish lighting. (By playing on cinematic conventions, it is possible to harness the expectations of an audience. See the section on Allusions to Existing Media, below.)

A dynamic musical score, featured in (void*), in particular, is also a very powerful tool for showing how the characters are feeling. By changing themes, tempo and timbre, the score can become a valuable asset for amplifying the emotional content of the scene.

6. Life Cycle

What does an embryonic synthetic character look (and act) like? When our group is developing a character, we come up with interim stages that will help us get a feeling for how the "adult" character will look and act. For example, when we were designing the raccoon for *Swamped!*, we wanted him to sneak around the barnyard. But what good is half of a sneaky character? (This is similar to another problem in evolutionary biology – "What good is half an eye?" the solution to which, Charles Darwin proposed, is that complexity arises by slow degrees[3].) Which is a more appropriate first stage of a character that sneaks around a scene: a character who walks around a scene, or a character who acts sneaky but can not navigate? This may very well depend on whether the focus of the installation is navigation or emotional expressivity.

With the advent of learning in our characters, however, this process is beginning to change its shape. The virtual dog, Duncan, who herds sheep in *sheep|dog*, was "born" without any knowledge of the voice commands that he would need to know to work well with his shepherd (the participant). We are gradually shifting from building characters as hard-coded adults to building characters that learn certain components of their behavior. This parallels the evolutionary move from animals with primarily hard-coded behavioral repertoires (e.g., insects) to those in which learning plays a significant role in the behavior of adult individuals (e.g., mammals).

7. Allusions to Existing Media

When a person approaches something new, they bring to it many expectations from other things that they have already experienced and that are familiar to them. This concept has been well explored, with regard to technology, in the work of Reeves and Nass [11]. People watch TV and movies, read books, go to art museums, play video and board games, know some science, and are hip to pop culture. Our characters and installations allude to a wide variety of media.

Many of these allusions we put there intentionally. These allusions frame the participant's expectations about how the interaction will proceed. In creating virtual characters, we consider what other characters our audience will know. Even the most homogeneous audience brings a wide variety of cultural references and expectations with them. Intentionally alluding to established conventions that they are familiar with can make participants feel welcome and comfortable interacting with the installation.

Swamped! had a variety of cartoon references, especially to Wile E. Coyote and the Road Runner. *(void*)* referenced Charlie Chaplin and Edward Hopper's painting "Nighthawks". *sand:stone* referenced "Ozymandias" by Percy Shelley. *sheep|dog* referenced a popular British sheep herding show called "One Man and His Dog", and several Guinness Beer commercials. (People often find allusions in our work that we did not consciously put there, as well.)

In *(void*)*, the basic interaction was borrowed from Charlie Chaplin. In his film, "The Gold Rush", Chaplin sticks two forks into a pair of dinner rolls and does a little dance with them, during which the rolls seem to be transformed into feet. Many people know this scene (or the scene in "Benny and Joon" where Johnny Depp does the same.) As soon as participants catch the allusion, they think about an entire suite of concepts that relate to that character – funny walks, canes, black-and-white movies, top hats. They know

it's okay to laugh at our characters, and are able to situate our scene in a cultural context.

Finally, the cinematography systems that we have written for the various installations utilize as many clichés as possible. Red under light signals the bad guy, close-ups show stronger emotion, and a shot of two characters demonstrate that there is some relationship between the two. By situating the characters in a cinematic setting based on the Hollywood style that we all know (and some of us love), people relax a little in the face of the unknown – artificial life, novel tangible interfaces, computers that interact with them.

8. Balanced Components

The design of the characters themselves creates expectations. In order for a character to work well, it needs to fulfill those expectations and be generally well balanced. This applies to its "mind" – emotions, motivations, and action selection – and to its "body" – model, animations, and motor system. We do not make human characters that look like humans because we cannot make human minds.

Interim stages of our characters often suffer from imbalance, as various sub-systems come on-line earlier, or dependencies in work flow make it impossible to polish one component until another is nearly done. For example, the models of the characters in *(void*)* were quite polished and smooth and walking around like automatons well before their minds had taken shape. By the time the installation was complete, however, their minds matched their bodies well.

The same notion of "impedance matching" components applies to installations as a whole. If any one character or other element sticks out too much (even for being very much more interesting than the rest of the characters) it will make the entire installation seem unbalanced.

9. From Characters To Installations

Over the last several sections, we have discussed methods to conceive of characters. This section proposed several paradigms by which these characters can combine into coherent and cohesive installations.

9.1 Installation as Story

An installation can be conceived of in the same way that we think about stories. A story is essential a scenario populated by characters that undergo emotional change. By creating characters that are able to undergo emotional change and situating them in scenarios, we are able to establish situations in which stories emerge from the interactions of the characters. Since our characters are still very simple (compared to real people or animals), the stories that emerge are also very simple. However, as virtual characters become more complex, we hope that the stories that emerge from their interactions will become more interesting. Further, by allowing the participant to affect the emotional state of one or more characters (which was the explicit focus of *(void*)*), we allow stories to emerge in which participants are *central* to the emotional arc of those stories. When conceiving of an installation as a story, we look at what emotional arcs our characters might undergo, and find effective mechanism by which to put the participant or participants at the heart of that character development.

9.2 Installation as Social Interaction

A group of virtual characters can be seen as a social situation. A participant is able to participate in the social interaction, either by assuming the role of one of the virtual characters or by being seen by the characters as another creature who just happens to be "outside the box". This kind of conception of installations is helping to establish a new kind of community that crosses the line between real and virtual intentional beings.

9.3 Installation as Single Interaction

All elements of an installation should focus on one central, simple interaction. Each component – characters, camera, music, interface, lighting, world, set design – should support that interaction. (This perspective is derived from the basic technique that Activision uses when designing its games.)

For example, in *(void*)*, we began with the decision to base the installation on the question "What would Charlie Chaplin do if he had access to modern computers?" This quickly led to our central interaction – the roll dance that he does in "The Gold Rush". Starting with a simple, unified interaction provides a central spine for the entire installation. Just as all the elements of a work of art should be unified aesthetically, all the parts of an installation need to support the central interaction.

10. Conclusion

In this paper, we have taken Disney's illusion of life as our inspiration and Dennett's intentional stance as our rationale toward the creation of autonomous animated characters. We have presented a variety of ways of thinking about how to build these creatures and how to integrate them into installations that show them off. We suggest that while great installations come ultimately from characters that can convey their beliefs, motivations and desires, this only can happen if the installation as a whole serves to stage, enhance and focus attention on those characters.

11. Acknowledgements

The ideas in this paper have been accumulated over the last several years as a result of a huge amount of work by all of the members of the Synthetic Characters Group and its affiliates. Thanks also to Madaline B. Tomlinson.

12. REFERENCES

[1] Blumberg, B. 1998. Swamped! Using plush toys to direct autonomous animated characters. *Proceedings of the conference on SIGGRAPH 98: conference abstracts and applications*, p. 109.

[2] Blumberg, B. 1999. (void*): A Cast of Characters. *Proceedings of the conference on SIGGRAPH 99: conference abstracts and applications*, p. 169.

[3] Darwin, C. 1999. *The Origin of Species*. Bantam Classic. (Originally published 1859)

[4] Dennett, D. 1987. *The Intentional Stance*. Cambridge, Mass.: MIT Press

[5] Ekman, P. 1982. *Emotion in the human face. 2nd Edition.* Cambridge University Press, Cambridge, UK.

[6] Johnson, M. P., A. Wilson, B. Blumberg, C. Kline, and A. Bobick. 1999. Sympathetic interfaces: using a plush toy to direct synthetic characters. *Proceedings of the CHI 99 conference on Human factors in computing systems.*

[7] Marinelli, D. and S. Stevens. 1998. Synthetic interviews: the art of creating a "Dyad" between humans and machine-based characters. *Proceedings of the sixth ACM international multimedia conference on Technologies for interactive movies.*

[8] Pandzic, I.S., T. K. Capin, E. Lee, N. Magnenat Thalmann, and D. Thalmann, A Flexible Architecture for Virtual Humans in Networked Collaborative Virtual Environments, *Computer Graphics Forum*, 16:3, 177-188, 1997.

[9] Papert, S. 1980. *Mindstorms*. New York: Basic Books.

[10] Perlin and A. Goldberg, Improv: A System for Scripting Interactive Actors in Virtual Worlds, Proceedings of SIGGRAPH '96, 205-216, 1996.

[11] Reeves, B. and C. Nass. 1996. *The Media Equation*. Cambridge: Cambridge University Press.

[12] Rickenberg, R. and B. Reeves. 2000. The Effects of Animated Characters on Anxiety, Task Performance, and Evaluations of User Interfaces. *Proceeding of the CHI 2000 conference on Human factors in computing systems.*

[13] Rose, C.F., Cohen, M., Bodenheimer, B., Verbs and Adverbs: Multidimensional Motion Interpolation - *IEEE Computer Graphics And Applications*, Volume 18, Number 5, 1999

[14] Russell, J. 1980. A circumplex model of affect. *Journal of Personality and Social Psychology*, 29:1161-1178.

[15] Thomas, F. and O. Johnson. 1981. *The Illusion of Life: Disney Animation.* New York: Hyperion.

[16] Tomlinson, B., B. Blumberg, and D. Nain. 2000. "Expressive Autonomous Cinematography for Interactive Virtual Environments." *Fourth International Conference on Autonomous Agents.*

[17] Tomlinson, B., M. Downie, A. Benbasat, J. Wahl, W. Stiehl. 1999. "sand:stone - Artist Statement." *Leonardo* Vol. 32, No. 5, p. 462-463

[18] Tweedie, L. 1997. Characterizing Interactive Externalizations. *Proceeding of the CHI 97 conference on Human factors in computing systems.*

[19] Yoon, S.Y., B. Blumberg, G. Schneider. 2000. Motivation Driven Learning for Interactive Synthetic Characters. *Fourth International Conference on Autonomous Agents.*

[20] Yoon, S.Y., R. C. Burke, B. M. Blumberg, G. E. Schneider. 2000. Interactive Training for Synthetic Characters, *AAAI 2000.*