Developing an Aesthetic: Character-Based Interactive Installations

Marc Downie, Bill Tomlinson, Bruce Blumberg
The Synthetic Characters Group
MIT Media Lab

In the Synthetic Characters Group at the MIT Media Lab, we make virtual characters and build interactive installations to showcase them. The research agenda of our group is to design ways of making computational characters by examining natural systems. All of our interactive installations, whether or not they feature entities who *look like* real animals, contain characters that are in some way derived from the behavior of real animals. Our installations provide a chance for people to meet our characters and to find out how they work. Over the last several years, we've made an assortment of installations about a range of topics — from terriers to cartoon chickens to romantic poems.

The process of making these installations has helped advance our research goals and has helped develop an equally important set of aesthetic principles for our diverse work. This article describes how we construct the aesthetics of our installations. We first explain the essential structure of our installation–building toolkit, dividing it loosely into intra- and extra-character components. Then we provide descriptions and images from each of the major installations we've created since our group came together in 1997. Finally, we discuss what we aim for as artists when we create a new installation.

A Synthetic Character Installation – Basic Anatomy

Our virtual characters are constructed from three main modules — a perception system, an action selection mechanism and a motor system [4]. The perception system identifies important states, contexts and objects in the world [8]. The action system works out what actions should be taken and how. The motor system handles the moment-to-moment choreography of performing these actions [7]. These three core modules are often augmented with additional systems – systems for handling emotional or motivational states, for managing sensory data from the virtual world or for navigating around the world. Our research has explored quite complex implementations with a particular emphasis on learning and adaptation. For example we've built perceptual systems that can recognize new contexts for performing actions, action systems that discover and test hypotheses about how to obtain goals and motor systems that can build new actions from pieces of existing animations.

Each character-based installation has a number of additional modules for supporting and presenting the characters. In particular, all the installations we have produced have included networked graphics systems for displaying windows into the characters' worlds, device management systems for handling participants' input from a wide variety of custom built devices [9], virtual cinematography systems [13] and sound output for giving voices to our characters [7]. While not the primary focus of our research, these supporting elements are essential to our work and essential to the aesthetic flexibility of our group. They too are maintained and shared across a number of installations.

Installation Chronology

In order to demonstrate the range of subjects and aesthetics that our systems can create we now present short descriptions of some of the major installations that we've built in the last five years.

Swamped!

Our first group installation, *Swamped!*, is an interactive experience in which an instrumented plush toy provides a tangible, iconic interface for directing autonomous animated characters [1]. By manipulating a stuffed animal (a bright yellow chicken) corresponding to a chicken onscreen, the participant influences how the virtual character behaves. In addition, a fully autonomous raccoon, hungry to eat the chicken's eggs, completes the barnyard scene. Each character has a distinct personality and decides in real time what it should do based on its perception of its environment, its motivational and emotional state and input from the participant. The characters incorporate a new model of behavior and emotion, a new multi-target motion interpolation and new techniques for real time graphics. Automatic camera

Figure 1: *Swamped!*

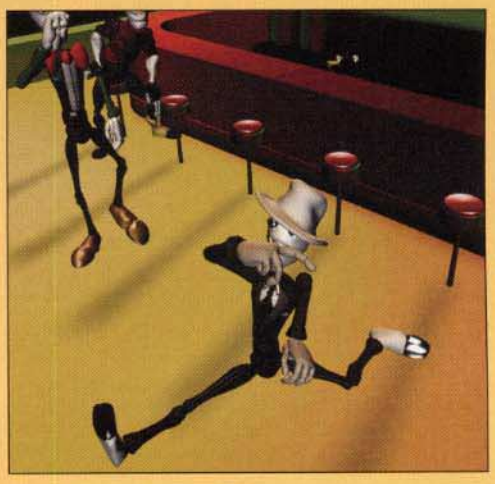

Figure 2: *(void*): A Cast of Characters*

Figure 3: *sand:stone*

Figure 4: *Duncan the Highland Terrier*

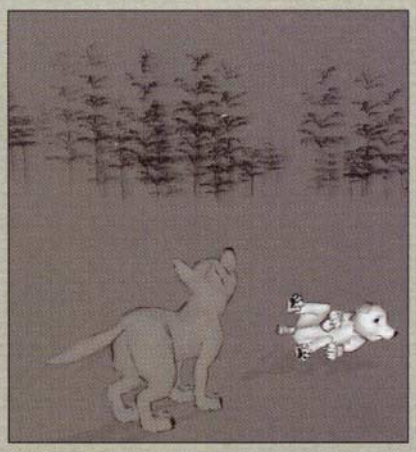

Figure 6: *AlphaWolf*

Figure 5: *music creatures*

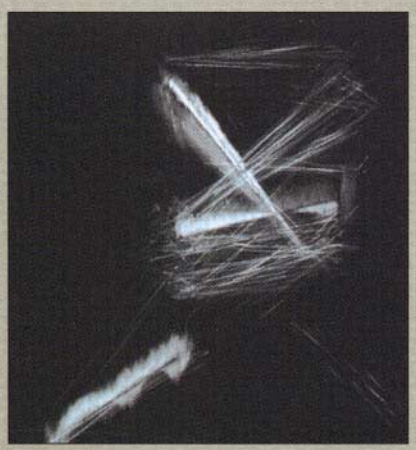

Figure 7: *Loops*

and lighting control help reveal the emotional content of each scene.

*(void *):A Cast Of Characters*

This installation, which premiered at SIGGRAPH 99, explores further the interactions between directable characters in a virtual space and users in the physical space [2]. The installation introduced a group of three characters sitting at a lonely bar. Through our "buns and forks" interface (inspired by Charlie Chaplin's film *The Gold Rush*), users could make the characters get up and begin to dance. It featured early work on preference and emotion learning in characters, a new interactive camera system and an interactive musical score generator.

These last two elements are key to the mood of the piece and tried to create a distinctively cinematic feel in an unpredictable environment. Our score generator started as an investigation into the question: how does one create a music score for an interactive piece? Such pieces have neither script nor central director, yet in *(void *)* we have a world and characters who seem to demand the full cinematic treatment. To facilitate the interplay between our characters and their supporting technologies, we built the music and camera using the same character toolkit that was used to build the creatures on screen.

sand:stone

sand:stone is a small piece inspired by Shelley's poem *Ozymandias* that explored another tangible interaction with characters but in a completely different mood [13]. The interface consisted of a collection of stones in a bed of sand; early prototypes of "musical creatures" were metaphorically located inside rocks. The placement of the stones in the sand was sensed by custom hardware, and affected the emotional state of an animated statue onscreen. The movement of the stones also had an impact on the interplay between the corresponding musical creatures, each singing with the other creatures through manipulated pre-composed musical material.

Duncan the Highland Terrier / sheep|dog: Trial by Eire

The creation of virtual dogs is a central research theme in our group [3]. In a series of installations featuring *Duncan the Highland Terrier*, we gave our virtual dog a variety of abilities. We built Duncans who could learn to perform actions in response to vocal cues, using traditional dog-training techniques; dogs that could herd (virtual) sheep in response to spoken commands; and versions that could be shaped or lured into performing novel actions.

Why dogs? Dogs force us to focus on certain hard problems in creating synthetic characters. For the participant, dogs represent a rich set of expectations and conventions that we need to meet. For us as designers, it presents a challenge that is at the bounds of plausibility (unlike, say, a virtual human) – and a challenge that is ideally suited for working on learning and expressiveness.

Dogs force us to focus on hard problems that we know have been solved (by nature) and that we know are worth solving. In addition, having dogs as a standard of comparison lets us know how well we are doing. The rest of our work, while it may not always look like a dog or sound like a dog, nevertheless benefits from these ideas.

music creatures

This installation is our first to concentrate solely on the idea of *music creatures* begun in *(void *)*. Here a networked colony of creatures plays very simple musical games. For example, creatures might swap musical material back and forth between themselves, or build up expectations of what they might hear next and dance accordingly.

Music creatures is, simultaneously, experiments in computer music and computer graphics — what people notice first about these creatures is not their playful murmurings but their bodies. These creatures possess bodies far less representational than those of our other creatures and often seem suspended between primitive geometric forms (lines and squares) and primitive organic forms (grasses, cells and trees).

The first *music creatures* installation was premiered at the opening of MediaLabEurope in Dublin, Ireland, in 2000. One of the creatures, a "listen" creature, was shown at SIGGRAPH 2000 as part of the *Trinkets: new pieces to play* [10] installation. A complete colony of creatures was shown in SIGGRAPH 2001's Art Gallery [7].

AlphaWolf

The *AlphaWolf* project premiered in the Emerging Technologies program of SIGGRAPH 2001 [14]. The installation features a pack of virtual wolves who form social relationships with each other. Three people interact with the virtual pack at the same time, each playing the role of one wolf pup in a new litter. Participants direct their pups with a microphone (into which they howl, growl, whine or bark to affect how their pups interact with their packmates) and a mouse(to tell their pups where to go).

Our wolves feature a simple model of social behavior incorporating learning, emotion, perception and development, and are rendered in a non-photorealistic, hand-drawn style.

Editor's Note: See this issue's VisFiles column for more information on *AlphaWolf*.

Loops

Our most recently exhibited installation is a collaboration that starts with motion capture data (taken by collaborators Paul Kaiser and Shelley Eshkar) of dance legend Merce Cunningham performing his piece for hands, entitled *Loops*. [5]. This animation material drives the movement of 42 small autonomous creatures that rebuild Cunningham's hands, making decisions concerning their appearance, the quality of their movement and their structural connections to other creatures. The work reuses and extends many of the technologies developed for the "music creatures," but sets them in a less interactive domain. Instead, this project focuses on the creation of a system (and an artistic process) that is complex enough to surprise us occasionally, but controllable enough to let us take advantage of those surprise discoveries.

Developing an Aesthetic — Some Unifying Goals

The installations described above clearly cover a large range of styles but, just as they share technological similarities, they also share a common set of aesthetic concerns. We have built these principles up through the practice of building installations; the result is a patchwork of principles that we can use in creating new installations. In this section we discuss some of these emergent aesthetic goals for our group's work, explain why we think that these goals are important and give some insight into how we can achieve them.

Stylistic Coherence

Despite our group's interest in building artificial intelligences grounded in animal behavior, our installations have never tried to present photo-realistically rendered, anatomically correct facsimiles of the real thing. Instead we have taken more indirect tacks – drawing on the traditions of cartoon, borrowing from slapstick or imitating black and white line drawings. It should be clear from the images above that this stylized approach has had an impact on the visible surface of our installation work. There are important reasons for these aesthetic constraints, reasons that underpin many of our design decisions.

Managing the expectations of novice participants is crucial to the success of a piece. Building a photo-realistic human, for example, would draw attention to the very things that we cannot do. For we cannot at present hope to satisfy the expectations that participants would have of such creatures – we would fail to create the same quality of animation as a live human, and we would certainly fail to create the same quality of intelligent behavior that such a style would suggest.

Instead, for example, *Swamped!* presents cartoon animal behavior, in a cartoon world, with cartoon anvils dropping from the sky and an interface straight out of childhood; *AlphaWolf* presents a consistent sketch of a coherent world – a stylized approximation of wolf interaction and social learning, a charcoal-like rendering style and an evocative interface. Similarly, *music creatures* do not possess human level musical capabilities; their bodies are not pianos, nor do they have hands or mouths. Rather, they are simple organic forms evoking a simpler biology while questioning preconceptions of what an "embodied intelligence" can look like.

Crucially, if some installation element does not contribute to this coherent world then it must be cut. For example, an early prototype of the *Swamped!* installation contained a beautiful and intricately modeled tree house located in the distance. Rather than adding to the installation, it detracted from the substance of the piece. Participants immediately wanted to be able to climb up into the tree house and when they could not, they were frustrated. The solution was to cut the tree house entirely – when it wasn't there, nobody could miss it. Therefore, our worlds tend towards simplicity; each screen of a music creature colony has only one creature floating in space, the world of *AlphaWolf* consists of nothing but a few very carefully rendered trees, and there are no buildings in *Swamped!* that are not interactive.

Whole behaviors can fall afoul of this "curse of content." An early, non-interactive prototype of *AlphaWolf* had the wolves simply wandering around, rendered to look like hand drawn charcoal. People passing by this prototype were capable of reading deep intentions into the images, which they found especially captivating. Maintaining this apparent depth while adding true behavioral complexity was in fact a significant challenge.

Finding an Interaction Metaphor: Animals

For an installation to be successful we must carry the stylistic coherence into the nature of the interaction itself. It is here that we have reaped the rewards of our focus on modeling animals. People know how to interact with animals and they enjoy doing it. In the unpredictable virtual worlds of an interactive installation we seek technologies that enable this particularly familiar and rewarding interaction. People play with animals; they empathize with animals; they try to work out why they perform certain actions. If we can persuade people that our animated renderings are creatures in this sense, and if we can build artificial intelligences that truly can be played with, motivated and understood – be they wolves or more abstract forms making sound – then people already know to interact with them.

The approach of the Synthetic Characters Group has been to tackle head on the problems of how to create such interactions by means of our research questions. What computational representations do characters need in order to have expectations? How are motivations incorporated into action selection? How should expressions of internal state be conveyed by a creature's motor system? The group is founded on the hypothesis that it might be more rewarding and easier in the longer term to engage these complex issues directly, rather than carefully scripting our way around them. *AlphaWolf*, for example, started from a desire to simulate the social behavior of wolves. Our work with the virtual dog Duncan stemmed from research into learning. These areas focus on the inner mechanisms of animals, rather than their surface appearances, and yet make our particular interaction aesthetics possible.

For example, the creatures in *(void *)* immediately look at their feet when the participant moves the interface rather than waiting for the gesture recognition system to complete its analysis. In *AlphaWolf*, the interaction begins with participants waking up their wolf pups, immediately helping form a connection between the humans and virtual creatures. Music creatures may habituate to repeated notes, or burst into enthusiastic activity after a period of boredom. One can tease them by building up a repeated musical texture only to stop short. While such musical creatures have no direct animal counterpart, being instead experiments in less representational graphics, we have balanced the extreme strangeness of their visuals with a reference to how we typically interact with animals.

Adaptive Installations

Building systems that are supposed to change while they continue to function is both a challenging engineering task and a challenging research goal. What are useful representations for a system that changes complexity? How should one go about authoring such a thing? Can we use learning as a tool for creating characters? There can be no better place to start on these problems than with animals, which are so well adapted for this very task.

It is not just our studies in animal behavior that compel us to build adaptation and learning into our creatures, it is also our desire to extend the lives of our characters and the complexities of the interactions that they can support. Each installation we have built has tended to include more and more learning as we have become better at addressing these issues and sketching solutions. *AlphaWolf*, for example, is built around social learning; the *music creatures* are driven by their perceptual learning abilities. Even in *Loops*, the creatures build up primitive expectations about their neighbors' movement that would be impossible to specify ahead of time, creating change on timescales far longer than the animation material used in the piece.

While machine learning is a huge field in its own right, such research is particularly important in an interactive installation. As we seek to make our installations interesting over longer timescales, our creatures *must* change over time. Once a character is "on screen" for more than a few minutes it will likely encounter a situation that it has seen before. To continue to appear intelligent that character must respond appropriately; this cannot be achieved without some element of adaptation on the part of the character. Building creatures with the ability to learn enables interactive works to be interesting on a second viewing, and the characters to remain convincing during sustained interactions in complex worlds.

Quality, Control

If there has been a wellspring of inspiration other than animals, it has been traditional cel animation. Animation has had an impact on how we view our work, how we build the technologies behind it and how we organize our creative process. We have, for example, built all our characters using "motor systems" that start from animation material created by animators, and then go on to blend, manipulate, layer that material rather than synthesizing motion from scratch.

Working with hand-crafted animations is by no means an easy option: techniques for building suitably flexible example-based motor systems remain an area of active research. Nevertheless we feel that starting here gives the best chance of capturing the inarticulable nuances present in the animator's craft and transferring the subtle cues that great animators use to create "the illusion of life" [11] into our creatures.

In order to allow our characters to determine the quality of their own motion, participants must not be permitted to control the fine motion of the characters. Even in our most representational work we have never created characters that can be directly puppeteered. Rather, we have developed interfaces that allow participants to have high level control over the actions of a character (i.e. *intentional control* [9]) without sacrificing the character's control over the quality of the motion.

Conclusions

None of the aesthetic goals we have described in this paper are easily achieved; all involve open research questions. What internal structures does a character need in order to have expectations, and how can it exhibit this behavior? How does one manipulate an animator's material without losing the animator's magic? How does one author a character that adapts and learns over time, and in what way should it change?

In addition, our research platform has provoked certain artistic experiments that extend beyond the scope of traditional artificial intelligence or traditional animation. Using our group's basic framework, our installations have explored new graphical vocabularies (*music creatures*), new ways of building collaborative art (*Loops*), new ways of staging interactive characters (*void* *) and many other avenues suggested by the very shared structures we have been building.

An interdependence between our aesthetic goals and our research should be clear. Our research agenda makes little sense without keeping our aesthetic targets in mind, and yet our goals keep provoking hard research questions. Fluidity between the design and implementation of large architectures and focused work on artistic experimentation is vital for the long term health of our medium as it moves beyond smaller, playful experimentation towards larger and inherently more collaborative works. Our synthetic character architecture and our set of installation-building tools represent our effort to build structures for a long term collaborative artistic practice.

References

1. Blumberg, B. "Swamped! Using plush toys to direct autonomous animated characters," *Proceedings of SIGGRAPH 98*: conference abstracts and applications, 1998.
2. Blumberg, B. "(void*): A Cast of Characters," *Proceedings of SIGGRAPH 99*: conference abstracts and applications, 1999.
3. Blumberg, B. "D-Learning: What learning in dogs tells us about building characters that learn what they ought to learn," *Distinguished Presentations in AI*, G. a. N. Lakemeyer, Bernhard, San Francisco, Morgan Kaufman, 2001, to appear.
4. Burke, R., D. Isla, M. Downie, Y. Ivanov and B. Blumberg. "CreatureSmarts: The Art and Architecture of a Virtual Brain," *Proceedings of the Game Developers Conference*, pp. 147-166, 2001.
5. Cunningham, M., M. Downie, S. Eshkar, P. Kaiser. *Loops*, Cambridge, MA, MIT Media Lab, 2001.
6. Downie, M. "Behavior, Animation and Music: The Music and Movement of Synthetic Characters," *Media Laboratory*, Cambridge, MA, MIT, 2001.
7. Downie, M. "music creatures," *Proceedings of SIGGRAPH 2001*: conference abstracts and applications, art gallery, 2001.
8. Isla, D. "The Virtual Hippocampus: Spatial Common Sense for Synthetic Creatures," *Media Arts & Sciences*, Cambridge, MA, MIT, 2001.
9. Johnson, M. P., A. Wilson, B. Blumberg, C. Kline, A. Bobick. "Sympathetic interfaces: using a plush toy to direct synthetic characters," *Proceedings of the CHI 99* conference on human factors in computing systems, 1999.
10. Paradiso, J., K. Hsiao, A. Benbasat. "Musical Trinkets: New Pieces to Play," *Proceedings of SIGGRAPH 2000*: conference abstracts and applications, 2000.
11. Thomas, F. and Ollie Johnson. *Disney Animation: The Illusion of Life*, New York, Abbeville Press, 1981.
12. Tomlinson, B., B. Blumberg, D. Nain. "Expressive Autonomous Cinematography for Interactive Virtual Environments," *Proceedings of 4th International Conference on Autonomous Agents*, 2000.
13. Tomlinson, B., M. Downie, A. Benbasat, J. Wahl, W. Stiehl, B. Blumberg. "sand:stone - Artist Statement," *Leonardo*, 32(5), pp. 462-463, 2000.
14. Tomlinson, B., M. Downie, M. Berlin, J. Gray, A. Wong, R. Burke, D. Isla, Y. Ivanov, M. P. Johnson, D. Lyons, J. Cochran, B. Yong, D. Stiehl, R. Soetjipto, D. Zaharopol, B. Blumberg. "AlphaWolf," *Proceedings of SIGGRAPH 2001*: conference abstracts and applications, 2001.

About the Contributors

Marc Downie, Bill Tomlinson, Bruce Blumberg
The Synthetic Characters Group, MIT Media Lab
Email: **{marcd, badger, bruce} @media.mit.edu**