Proceedings of the 2002 ACM SIGGRAPH Symposium
on Computer Animation. San Antonio, TX.

Leashing the AlphaWolves: Mixing User Direction with Autonomous Emotion in a Pack of Semi-Autonomous Virtual Characters

Bill Tomlinson, Marc Downie, Matt Berlin, Jesse Gray, Derek Lyons, Jennie Cochran, and Bruce Blumberg
Synthetic Characters Group, The Media Lab, MIT
{badger | marcd | mattb | jg | dlyons | jenelec | bruce}@media.mit.edu

Abstract

We present a system in which computer-graphical virtual characters may be controlled by a user and also remain "in character." The system allows the user to have high-level control over the actions of a character, while the emotional state of the character is autonomously maintained by the computer. We show how this system functioned as part of the *AlphaWolf* installation, presented in the Emerging Technologies program at SIGGRAPH 2001. Results from a 32-subject human user study support the hypothesis that users could control a character's actions without sacrificing its realistic autonomous personality. This system is appropriate to the control of computer-graphical entities that are meant to have personalities distinct from those of the humans that direct them.

CR Categories and subject descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism --- Animation, Virtual Reality; K.8.0 [Personal Computing]: General --- Games; H.5.2 [Information Interfaces and Presentation]: User Interfaces --- Interaction Styles.

Additional Key Words: Directable Characters, Autonomous Agents, Human Factors, Emotion

1 Introduction

The entertainment industry has many compelling characters – Bugs Bunny, Mickey Mouse, Buzz Lightyear, Shrek. These characters are incredibly powerful in the linear media of film and television. However, making interactive versions of these wonderful characters is a hard problem. As soon as a person has control over the behavior of a character, there is the strong possibility that the person will make it do something inappropriate. How can entities be controlled by a user and yet stay "in character"?

The *AlphaWolf* installation (see Figure 1 and the Video Figure), which premiered in the Emerging Technologies program at SIGGRAPH 2001, offers a step towards a solution to the problem described above. This interactive installation features a pack of three-dimensional animated wolves whom participants can direct, and whose behavior is based on the natural behavior of the gray wolf (*Canis lupus*). The challenge of building directable virtual wolves who nevertheless exhibit plausible wolf behavior is similar to that of preserving a pre-existing personality in an interactive character. Just as there is usually a clear answer to "What would Bugs Bunny do?" in a given situation, there is a clear answer to "What would a real wolf do?" While we are not claiming to present a solution to the full "Bugs Bunny Problem", we hope that the approach presented here might help inform further efforts in this domain.

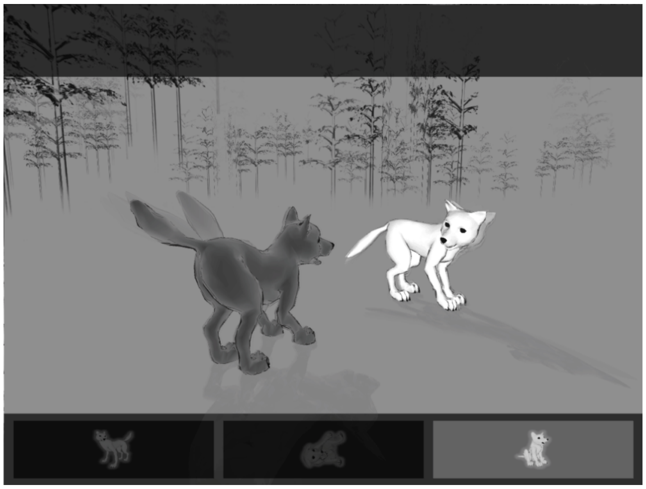

Figure 1: Two virtual wolf pups interacting in the *AlphaWolf* installation at SIGGRAPH 2001.

In *AlphaWolf*, each human participant has high-level control over the actions of a virtual wolf, but the wolf's emotional state is autonomously determined by a computer. This distinction between action and emotion makes it possible for participants to direct the actions of the virtual wolves without compromising the wolves' realistic behavior. Since participants retain the ability to direct their characters, they can become immersed in the experience of "playing" the wolves. At the same time, the autonomously determined emotional states cause the wolves to behave like real wolves throughout the course of the installation. In the absence of user control, the wolves are capable of choosing socially appropriate actions and responses. Nevertheless, the virtual wolves always do what they're told and are expressive enough to convey how they feel about the courses of actions chosen for them. We propose that a clear division between action and emotion is a useful mechanism for making semi-autonomous characters who obey the direction of a human participant and still present a consistent personality.

In the next section, we describe the *AlphaWolf* installation as presented at SIGGRAPH. We then describe a selection of related works that are most relevant to this system. We next turn to the mechanism itself, exploring the action/emotion distinction in greater depth. Then we present the results of a human user study on participants in the *AlphaWolf* installation. Finally, we summarize this paper's claims about why combining user-directed action and autonomous emotion is a useful technique for creating semi-autonomous directable characters.

2 The *AlphaWolf* Installation

The *AlphaWolf* installation premiered in the Emerging Technologies program at SIGGRAPH 2001. The installation features a pack of directable three-dimensional animated wolves. Three participants direct the actions of three pups in a newborn litter. By howling, growling, whining or barking into microphones, participants tell their pups to howl, growl, whine or bark (see Figure 2). In addition, by clicking with a mouse in the virtual world, participants can tell their pups where to go and with whom to interact.

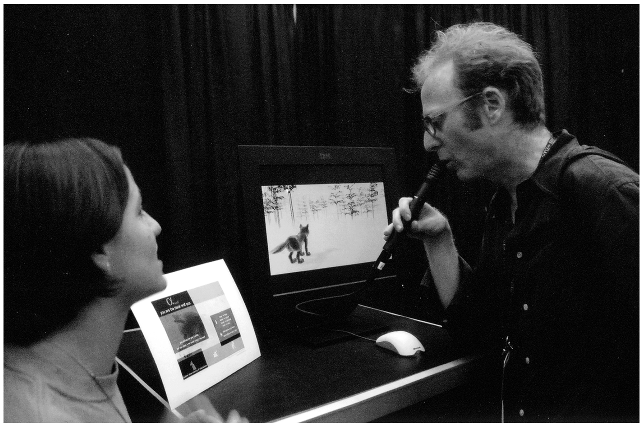

Figure 2: A participant directs the actions of his wolf.

An interaction session with the wolves lasts approximately five minutes. During this time, the pups wake up and meet their pack mates. There are six wolves in all – three user-directed pups, and three fully autonomous adults. The individual wolves autonomously form and remember social relationships with each other based on their interactions with other wolves. In turn, these relationships color the way in which the pups perform the actions that they are directed to take.

The wolves are rendered using a custom-written "charcoal renderer." The renderer uses a technique based on programmable vertex shaders to give the wolves their characteristic look. Through the rendering style, character design, animation, sound design and cinematography, we tried to capture the desolate feel of the arctic tundra.

2.1 Interface

As described above, people interact with the AlphaWolves through two main interfaces – a microphone and a mouse. Each participant's microphone is backed by a system that performs acoustic pattern matching on the utterances that it receives. We use a simple mechanism for classifying sounds involving utterance length and harmonicity.[1] This system proved to be remarkably effective at capturing the distinctions between the four utterance types that *AlphaWolf* uses, especially in the noisy environment of the SIGGRAPH floor. When a user clicks on the screen with the mouse, the wolf moves to that point. If the participant clicks on another wolf, that wolf becomes the target to which the pup runs.

In addition, each time a wolf pup meets a member of its pack, a button appears on the border of the user's screen. This button has an image of the other wolf on it. If a user clicks on one of the buttons, the pup will run to the last place where the pup saw that wolf. The button reflects the pup's set of beliefs about that other wolf. The icon on the button changes to demonstrate the relationship that the pup has with that wolf. For example, if another pup has habitually submitted to the user's pup, the button for that pup will show it in a submissive pose. Similarly, if the other pup habitually dominates the user's pup, the button will have the other pup looking dominant. The buttons allow the user to "see into the pup's mind," and understand the pup's relationships better. An important element of the buttons is that they do not offer any information that is not represented inside the virtual brain of the wolf pup.

2.2 Example Interaction

As an example of the kind of interaction that occurs in *AlphaWolf*, imagine that a participant growls into the microphone to direct his gray pup to growl at the white pup. In accordance with his user's direction, the gray pup takes on a dominant stance and growls. Seeing the gray pup growling, the white pup's user directs his pup to whine. The white pup duly rolls over on his side, whining. Because of this interaction, each pup forms a social relationship with the other; the next time they meet, their remembered relationships will affect the way in which they interact. If directed by his user to interact with the gray pup, the white pup might approach him submissively, and might cringe if his user directs him to growl. Despite his cringe, the pup will perform the action his user directs, but in a style that demonstrates the pup's own impression of his relationship with the gray pup.

2.3 Relationships

The essence of the *AlphaWolf* model of social relationship formation involves emotion, perception, learning and development. Each wolf maintains an emotional state that is affected by its interactions with the world. A wolf is able to perceive the identity of its pack mates, recognizing them as distinct individuals. It forms an emotional memory of each individual after its first interaction with it. When it again encounters that individual, the emotional memory influences its current emotional state, so that it can "pick up where it left off" with regard to its emotional relationship. At the end of each interaction, it revises its remembered emotional relationship with that social partner. This mechanism is most closely modeled after Damasio's Somatic Marker Hypothesis [1994]. This mechanism is described more fully in [Tomlinson 2002 (to appear)].

Because of their ability to form social relationships autonomously, the pups are not simply puppets; even though human participants are directing their actions, they develop their own personalities and display them over the course of their young lives.

One of the benefits of the emotional memory mechanism is that it keeps track of the pups' relationships so that participants don't have to. Novice users often find it challenging to remember their pup's relationships with each new wolf; the icons on the buttons help remind them of their interaction history.

---

[1] Howls are long and harmonic. Whines are short and harmonic. Growls are long and non-harmonic. Barks are short and non-harmonic.

2.4 Waking up

Many participants were a little uneasy at first about the prospect of howling into a microphone in front of a crowd. We had noticed this phenomenon in some preliminary runs of *AlphaWolf*. To ease people into the use of the microphone, we caused all three pups to start off asleep, and to be woken up by any noise made into the microphone. By the time most people had awakened their pup and engaged in their first interaction with another wolf, they appeared to be immersed in the interaction and unconcerned about making all kinds of wolf noises in front of a crowd. This transformation from anxious observer to enthusiastic participant causes us to believe that directing a pup in the *AlphaWolf* installation provided participants with a compelling interaction.

In addition to putting users at ease, the waking-up of the pups served to shape the relationships between users and pups. Semi-autonomous characters such as the virtual wolf pups are somewhere in between avatars and autonomous agents. Throughout the installation, the distinction between user and pup is clear, but there is still a strong empathetic association between the two. The waking-up process helps to focus this association by means of an intimate camera angle, a snoring sound effect, and an ear-flip action that the sleeping pup does when he hears a sound. Installation elements that help reinforce the relationship between users and their characters are invaluable to the process of immersing users in the interactive experience.

3 Related Work

This paper's system for controlling semi-autonomous characters is derived from a variety of sources – virtual characters and agents, computer games, computational models of emotion, and the behavior of wild gray wolves. While this research is inspired by many of the projects below, and incorporates elements of several, we are not aware of any previous project that has made a distinction between user-controlled action and autonomous emotion.

3.1 Virtual Characters and Agents

Various researchers have developed mechanisms for the high-level direction of virtual characters, often drawing inspiration from existing art forms such as improvisational theater. Hayes-Roth *et al.* [1995] have explored "directed improvisation" as a way for users to direct and constrain the behavior of computer characters. Johnson *et al.* [1999] discuss the notion of "intentional control" – interpreting user input to allow the user to control a character at the behavioral level rather than at the motor level. Blumberg and Galyean [1995] integrated autonomy with directability using a multi-level approach. The Improv system of Perlin and Goldberg [1996] addresses the creation of believable synthetic actors, using procedural techniques to create layered, non-repetitive motions and transitions. Johnson [1994] created a system for creating and testing semi-autonomous animated characters. Assanie [Assanie 2002] offers a system for integrating directable characters with a centralized narrative manager.

Other research projects have focused on different elements of creating believable characters. Cassell, Badler and others (*e.g.*, [Cassell 1994; Cassell 1999; Badler 2001]) have explored making conversational characters who express themselves through voice and gesture in lifelike ways. Thalmann, Magnenat-Thalmann and others have also been working on virtual humans, particularly to serve as virtual actors (*e.g.*, [Thalmann *et al.* 1997; Magnenat-Thalmann *et al.* 1998]). Bates and his colleagues (*e.g.*, [Bates *et al.* 1992; Reilly 1996]) and their company Zoesis have done research on making virtual characters with expressiveness, emotions, and social behavior to serve in interactive story environments. Blumberg and his colleagues have built autonomous and semi-autonomous synthetic characters with learning, emotion and social behavior (*e.g.*, [Blumberg 1996; Burke *et al.* 2001]).

Various researchers have explored ways of expressively controlling the bodies and faces of animated characters (*e.g.*, [Hodgins and Pollard 1997; Brand 1999; Chi *et al.* 2000]). Rose's research [1999], for example, describes a motor control system with an explicit separation between the action itself and the style of the action. The motor control system that we use in *AlphaWolf* reflects Rose's verb/adverb distinction, which parallels the action/emotion split.

The Autonomous Agents community addresses similar problems to those that we confronted in making *AlphaWolf*. The problem of making an autonomous creature that can be controlled by people has been a long-standing challenge for Agents researchers (*e.g.*, [Strassman 1994]). In the Autonomous Agents 2001 conference, for example, there were several papers addressing the problem of autonomy and user-control (*e.g.*, [Scerri 2001]). Our work is distinguished from these projects by a focus on graphical expressiveness and inspiration from natural systems.

The spectrum of virtual creatures stretches from user-controlled digital puppets to fully autonomous agents. User-control and autonomy can be combined to make virtual entities with elements of both kinds of control. The researchers above have explored various means of striking this balance. While the goal of the system described here is similar to the goals of the projects described above, this system's approach to the problem is novel.

3.2 Computer Games

The challenge of making directable characters with strong personalities is relevant in a very practical way in the making of computer games. A number of games share certain elements with the *AlphaWolf* installation. Maxis' *The Sims* features semi-autonomous characters who can be directed to take a variety of actions. These actions in turn affect their emotional states and relationships with each other. The success of *The Sims* shows that social and emotional phenomena can be quite engrossing to game players. In Lionhead's *Black & White*, the player's autonomous creature changes its shape and the style of its behavior to reflect the user's interactions. *Rockett's New School* by Purple Moon explores another angle on game-play – players are asked to choose the emotional style in which their character should respond to events, and the character autonomously chooses behaviors in accordance with that emotional state. This approach is nearly the inverse of the mechanism described here, but explores a similar action/emotion split. *AlphaWolf*'s focus on an autonomously developing and continuously changing social life, grounded in clearly expressed emotional experiences, separates the work presented in this paper from the characters traditionally found in computer games.

3.3 Emotional Models

In order to give the virtual wolves an emotional state to serve as the basis for their autonomous relationships, it was necessary to choose a computational representation that captures the necessary range of emotional phenomena. Much research has already been done both in understanding emotions and in simulating them computationally. Darwin's ideas about emotions [Darwin 1965 (originally published 1872)] form the basis for much of modern research into understanding emotions scientifically. For a far more comprehensive discussion of emotional models in computational systems, the reader is directed to Rosalind Picard's book, *Affective Computing* [Picard 1997].

For the *AlphaWolf* installation, we considered two main emotional models – a categorical approach and a dimensional approach – to represent the wolves' emotional state. The categorical approach separates emotional phenomena into a set of basic emotions. Ekman's model, for example, categorizes the range of emotions into fear, anger, sadness, happiness, disgust and surprise [Ekman 1992]. This model provided the basis for an implementation by Velasquez [Velasquez 1998]; others (*e.g.*, [Gadanho and Hallam 1998]) have also implemented categorical models.

The dimensional approach (*e.g.*, [Schlosberg 1954; Smith 1989; Plutchik 1991; Russell 1997]) maps a range of emotional phenomena onto an explicitly dimensioned space. Various researchers have implemented versions of the dimensional approach; for example, Breazeal [Breazeal 2000] used a three-dimensional space (Arousal, Valence, Stance) to give affective tags to occurrences perceived by her robot, Kismet. These tagged events in turn affected the emotional state of the robot.

We found the dimensional approach to capture more effectively the range of behaviors exhibited by the gray wolf.

3.4 The Gray Wolf

Using examples from nature is a well-established tradition in the building of virtual creatures (*e.g.*, [Reynolds 1987; Terzopoulos 1994; Allison 1996]). We have chosen the gray wolf (*Canis lupus*) as the specific natural model for this simulation for several reasons. First, wolves manifest distinct social phenomena that are complex enough to be interesting, yet clear enough to provide direction for the simulation. Second, because of televised nature programs and other sources, many users have some foreknowledge about how wolves behave. Finally, the social behaviors of wolves are similar enough to those of humans that the mechanism in this paper may be relevant to human social behavior and simulation. In addition, having a concrete model from nature has provided a steady direction for this research.

In their natural environment, gray wolves form complex social groups known as packs. The core of most packs is a family – a breeding pair of adults, their puppies, and sometimes a few adult offspring of the breeding pair [Murie 1944; Mech *et al.* 1998]. The average pack size is approximately 7-9 individuals, but some packs may contain more than 25 wolves. Wolves maintain their social relationships through ritualized dominance and submission behaviors. There are two main types of submission that wolves exhibit – passive submission and active submission. Passive submission involves a wolf lying on his side or back, exposing the ventral side of his chest. The ears are held close to the head, and the tail is tucked between the legs. Active submission involves a crouched posture with backward directed ears, and licking or pecking the mouth of the dominant wolf [Schenkel 1967; Fox 1971]. Dominance behavior involves a "confident" posture with tail and ears erect, and direct staring at the social partner. More vigorous dominance behaviors involve pinning the muzzle of the submissive wolf to the ground and growling [Schenkel 1967].

4 Mechanism

As we have described above, action and emotion are controlled separately in the *AlphaWolf* installation – action by a human participant, and emotion by an autonomous control system. In order for action and emotion to be controlled separately, it is necessary for the wolves' "brains" to have a clear separation between the two. In order for participants to have control over their pups' actions, the pups must have an internal representation of action that is amenable to control. Similarly, for emotion to play a significant role in how pups perform actions, there must be an internal representation of emotion, and a mechanism by which it affects the style of actions. This section presents action and emotion in turn, and then addresses the interplay between the two.

4.1 Action

In *AlphaWolf*, a participant has direct control over the actions of his pup. Two elements in this system make this control possible: a clear computational representation of actions, and a way for the participant's interface to influence the selection of those actions.

Each virtual wolf is able to perform an assortment of actions – for example, sleep, stand, walk, dominate, submit, or howl. These actions are discrete elements within the wolf's behavior system. The representation of action that we use is derived from Burke *et al.* [2001]. Each action comes bundled with a trigger context, which determines when the action happens, and a "do-until" context, which determines when the action finishes. An action may also have an object, which is the target of the action. For example, a sleep action's trigger context might be "when fatigue is above a threshold", its do-until context might be "when fatigue drops below some other threshold", and its object might be "near the den". This action competes with other actions based on the values of their respective triggers. Once an action becomes active, it stays active until either its do-until is satisfied or some element of the world changes significantly (*e.g.*, "Someone just growled at me.")

This representation of action makes it easy to incorporate user-control. When the acoustic pattern matching system associated with a certain microphone recognizes a specific utterance, it feeds a value into the trigger context of the appropriate action. For example, if the user howls, the pattern-matcher tells the howl action's trigger to go high. It is possible to blend a user's input with autonomous control – both can contribute to the trigger values of the actions.

The control that this system gives a user is at a fairly high level – the level of an individual action. Most people are not good puppeteers; rather than being asked to control every joint angle in real time, users are allowed to influence the behavior system at the "action" level. Users direct their pups at a level that seems natural, and one that is mirrored by the internal structure of the behavior system.

Causing participants to direct their pups at this high level has several benefits. First, it gives them a strong sense of control because they influence the pups at the same level that people use when we think about actions. People don't often think about a periodic, cyclical bending of the knees, hips and ankles; we think about "walking." Directing a pup at the level of "going over there", "howling", or "whining" causes users to *perceive* themselves as having a high degree of control. Second, it makes interacting easier for them, because they do not have to struggle with real time control over the 39 rotational joints in each wolf.

Finally, it allows the emotion system (see below) to have an impact at a lower level of control, where the participants' inputs get translated into joint rotations. Users feel like they have complete control because they do not think about the low-level control while they are howling, growling, whining and barking.

4.2 Emotion

As we described above, participants have direct control over the actions of their pups; how those actions are executed is determined autonomously. The emotion system plays a significant role in the conversion of user direction to motor action.[2] The two central elements of the emotion system are: a computational representation of emotion, and a mechanism by which that emotion influences the style of actions that the pup performs.

The central representation of emotion in the *AlphaWolf* system is quite simple – a single floating-point value of *dominance*, which varies from 0.0 to 1.0. Each wolf's dominance value is affected by his interactions. For example, being growled at causes a wolf's dominance to drop, and being the target of another wolf's submission causes a wolf's dominance to increase. We decided to make the emotion model very simple so that it would be clear to participants at SIGGRAPH, who only interacted with the installation for approximately five minutes each.

The emotional state of the wolf at a given time is affected by several factors: the wolf's previous emotional state, an innate drift, the presence of certain innate releasers in its perceptual environment (*e.g.*, being growled at reduces dominance), and the presence of learned releasers (*e.g.*, the black pup is here, who usually dominates me, so my dominance will be reduced by his presence).

In order to give the wolves a continuous range of expressive behavior, our animator crafted example animations at the extremes of the range. These extremes are blended together in real time by the wolves' motor system. The dominance axis defines the primary range of expressiveness available to the wolves. Other expressive, blended axes are used for navigation (left vs. right) and for aging (young vs. old).

Through the expressive range of dominance, the pup's emotion system affects the style with which the pup performs the actions suggested by the user. For example, if a pup is asked to growl at a more dominant wolf, that pup will growl in the most submissive way possible. Rather than standing tall, raising its tail, and holding its ears erect, the pup will crouch, tuck its tail between its legs, and flatten its ears against its skull while growling. This behavior is still recognizable as a growl, but gives a very different impression about the motivational state of the wolf than would the more dominant version of the same action.

Allowing the wolves to determine their own emotional state makes them appear to have personalities, rather than just being puppets. If, when the user directs the pup to growl, the pup immediately cringes, there is the strong impression that the pup has a distinct opinion about that action. The pup clearly has its own beliefs and desires.

4.3 Interplay

There is some essential interplay between action and emotion in semi-autonomous characters. As Antonio Damasio [1994] points out, emotion is central to the way people (and animals) decide what to do. How can we make a user responsible for the actions of a character, and yet allow emotion to play a significant role?

This system's mechanism for causing a user to behave in a manner consistent with the desires of his virtual wolf is to make the interaction easier, faster or more fluid if the user's actions match the wolf's desires. For example, when a user asks his pup to growl at a submissive wolf, the pup runs to do so. If the user asks his pup to growl at a dominant wolf, on the other hand, the pup will walk forlornly. Since running is faster than walking, the user is rewarded for having his wolf behave "in character" by arriving at a destination more quickly.

In order to determine if a user's actions match the wolf's desires, it is necessary for the wolf to *have* desires. The behavior system of each wolf is able to engage in autonomous behavior and enables the wolf to determine autonomously whether the user's suggestion matches what the wolf would "naturally" do. The pup then uses this match to perform some simple low-level action control, *e.g.*, choosing a gait.

If a user does not interact for a period of time (~15 seconds), the pup will begin to behave autonomously, interacting with its pack mates in ways appropriate to the relationships it has formed. However, as soon as user input resumes, it overrides the autonomous behavior.

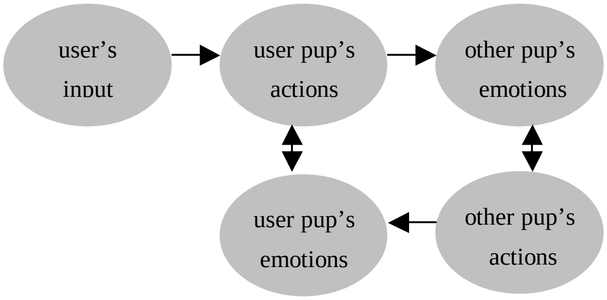

Figure 3: The flow of control among a user and two pups.

Certain behavioral elements are under the control of both user and autonomous system. The attention mechanism, for example, which determines where a wolf looks, has elements of both kinds of control. When a user clicks on the button for another wolf, to tell his pup to interact with that individual, the pup will look over at that wolf. However, among real wolves, submissive individuals rarely hold eye contact with dominant individuals. Therefore, the pup may look away, occasionally glancing back at the dominant individual. In addition, pups autonomously react to interactions initiated by other wolves by switching their object of attention.

Figure 3 summarizes the way in which the actions and emotions of the semi-autonomous pups interact. In this figure, we see that the user's input affects the actions of his pup, which in turn affect the emotions of another pup. The emotions of that pup help determine the actions that it will take. That pup's actions will affect the emotional state of the user's pup, which in turn colors how the pup takes the actions that the user directs. While other factors may affect the pups' actions and emotions (*e.g.*, the other

---

[2] Several other systems also play a role in the interpretation of user control. The navigation system, for example, determines how much *run_left*, *run_straight*, and *run_right* a pup should use to orient and reach the target when a user clicks. Also, if a user tells a wolf to howl, the wolf will autonomously decide whether to howl from a sitting, lying or standing position depending on which pose is closest to its current configuration.

pup may be user-controlled as well, or may have autonomous drives besides the emotional state depicted here), this diagram describes the essential path by which a user's control affects various elements of the wolves' behavior.

Here is a more concrete example of the various elements discussed above. A user directs his pup to submit to a certain other pup on the first time the two pups meet. His pup does so. The submission of the user's pup affects the emotional state of the other pup, making it feel more dominant. Because it is now feeling dominant, it exhibits dominant behavior toward the user's pup. Being dominated causes the user's pup to learn a submissive relationship toward the other pup. On the next occasion when the two pups meet, the user directs his pup to growl at the other pup. The user's pup duly growls but, remembering his submissive relationship to the other pup, does so in a crouched posture, looking away as he growls. Throughout this process, the user has been in control of his pup's actions at a high level, but the pup has developed and exhibited its own personality and relationships.

The *AlphaWolf* installation features virtual wolves who are both directable and plausibly wolf-like; these two components are made possible by the division of control between action and emotion. This division allows users to become immersed in interacting with their pups, without sacrificing the developing personalities of the pups that emerge over the course of the interaction. While this personality depends to a great extent on the kinds of interactions that users cause their pups to engage in, it becomes less dependent on the user as the interaction proceeds.

5 Evaluation

The essential claim of this paper is that the *AlphaWolf* system gives people control over the actions of their characters without sacrificing the realistic behavior of those characters. The degree of control granted to participants might have jeopardized the realism of the behavior because it allowed participants to direct their wolves to perform socially inappropriate actions. For example, a user could make his wolf growl at a social partner who is clearly dominant to him, a behavior that a real wolf would not habitually perform. Despite this control, participants appeared to continue to believe in the realism of the virtual wolves. In order to examine whether or not our system in fact preserved the wolves' behavioral realism while allowing users to have control, we performed a human user study involving 32 subjects.

The subjects ranged in age from 17 to 55 (mean = 26.2, std. dev. = 7.8). Half (16) were female and half were male. Each subject began with a short (2.5 minute) clip of a National Geographic Video about a pack of arctic wolves [Rosenfield 1988]. After the video, each subject watched or interacted with three four-minute runs of virtual wolves. Each run featured three wolf pups – one gray, one black and one white. The pups started each run with no emotional relationships. In two of the three runs, the subject was assigned to direct the gray pup to form specific dominant or submissive relationships with its siblings. In the remaining run, the subject did not interact with the wolves, instead watching the pups interact with each other autonomously. The three runs were administered in random order. After each run, the user was asked to fill out a questionnaire ranking his or her opinion on a range of topics, using a Likert scale from 1 to 7.

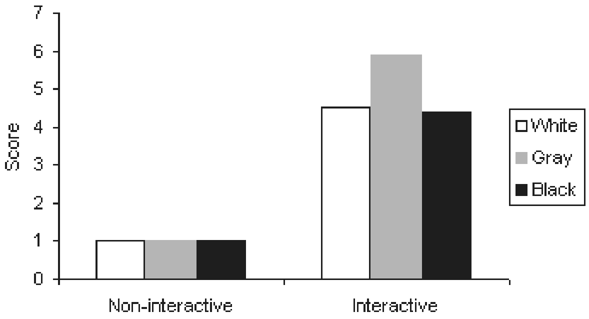

Figure 4: Subjects felt significantly more control over the behavior of all three pups in the interactive runs.

The two question topics relevant to this paper were the amount of (direct or indirect) control that the subjects felt they had over the behavior of the various pups, and the similarity of the virtual wolves' social behavior to the social behavior exhibited by the wolves in the National Geographic Video.

Not surprisingly, subjects felt significantly more control ($p < 0.0001$[3]) over the pups' social relationships during the interactive runs (n = 64, average score = 4.8), than during the non-interactive runs (n = 32, average score = 1.0). Subjects also preferred interactive to non-interactive runs with regard to their control over the behavior of the white pup (4.5 vs. 1.0, $p < 0.0001$), over the behavior of the gray pup (5.9 vs. 1.0, $p < 0.0001$), and over the behavior of the black pup (4.4 vs. 1.0, $p < 0.0001$) (see Figure 4).

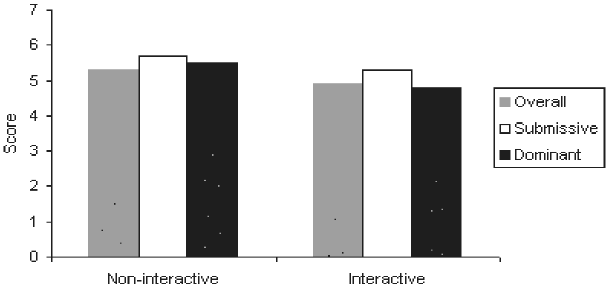

Figure 5: The social behavior of the wolves was almost as realistic in interactive runs as in non-interactive runs.

The fact that people felt more control over the gray pup than over black or white is not surprising either, since that was the pup over whom they had direct control. Nevertheless, the fact that they also felt significant control over black and white suggests that they were successful in having an indirect impact on those pups through their direct control over gray.

Not only did subjects *feel* that they had control, but they in fact *did* have control. Recall that each subject was assigned at the beginning of each interactive run to direct the gray pup to form specific relationships with black and with white. In 93.8% of these relationships, the internal representation of the gray pup's relationship at the end of the run matched the relationship that had been assigned to the subject at the beginning of the run ($p < 0.0001$[4]). Chance would have predicted 50% success at this task.

---

[3] P-values in this section were determined using the Binomial test.

[4] The Mann-Whitney U test was used to calculate this p-value.

The strong correlation between the assigned relationships and the relationships that the pups had actually formed demonstrates that subjects had succeeded in directing their pups to form specific relationships.

The two results described above, one involving users' subjective experience of control and the other involving their performance on an objective task, confirm the first part of this paper's hypothesis, that the *AlphaWolf* system gives people control over the characters that they are directing.

The second part of the hypothesis that we sought to confirm is that during the process of controlling an entity, the entity nevertheless maintains its realistic behavior or "personality". To test this part of the hypothesis, we included questions on the questionnaire about how similar the wolves' social behavior was to that of the real wolves in the video.

The results of this set of questions demonstrate that subjects found the virtual wolves to be almost as realistic in interactive runs as they were in non-interactive runs. The interactive runs ($n = 64$) received only slightly lower scores than the non-interactive runs ($n = 32$) on "similarity of overall social behavior" (5.3 vs. 4.9 on a 7 point Likert scale, $p = 0.2$[5]), "similarity of submissive behavior" (5.7 vs. 5.3, $p = 0.2$) and "similarity of dominance behavior" (5.5 vs. 4.8, $p = 0.01$) (see Figure 5).

The figures above demonstrate that being able to control a virtual wolf caused only a slight reduction in the realism of the wolves' behavior. The reduction did have statistical significance in one of its three parts ("similarity of dominance behavior"); however, the amount of reduction in all three cases was relatively minor, as is clearly visible in Figure 5, while the increase in control (visible in Figure 4) was striking. This result supports the hypothesis that people can control the actions of a character without dramatically compromising that character's realistic behavior.

The experiences that people had interacting with the *AlphaWolf* installation at SIGGRAPH 2001 support these results. While we were unable to collect statistical information, the great majority of the 500-1000 people who interacted with the installation appeared to become immersed in the experience and to feel successful in controlling their pups.

6 Conclusion

When we began to develop the *AlphaWolf* installation, we wanted to make an installation that featured a pack of virtual wolves who act like real wolves. In addition, we wanted to give people the ability to control those wolves without sacrificing their realistic behavior. We developed the system in this paper to meet these two goals.

The *AlphaWolf* installation allows a user to have high-level control over the actions of a character. The character itself retains autonomous control over its emotional state. This separation of control over action and emotion creates an interactive experience that is both controllable and true to the natural model on which it is based.

The mechanism for semi-autonomous characters described in this paper has ramifications in a variety of aspects of interactive installation design. The character architecture needs to integrate cleanly with the action/emotion division. The interface design needs to give users control of a character's actions while withholding control of its emotional state. Care must be taken to craft the relationship between the user and the character, who is neither avatar nor fully autonomous agent.

There are a number of ways in which the system presented here could be extended. For example, the system could incorporate a more elaborate emotional model, or some other means of increasing the behavioral complexity of the characters. Alternatively, the system could be extended to work with different paradigms for directing the characters. We believe that the system described in this paper is simple enough and general enough to accommodate a wide range of future work.

The problem of making a wolf who is both interactively controllable and plausibly wolf-like is similar to a problem faced by the entertainment industry, who wish to make virtual characters with strong personalities that are nevertheless controllable. We hope that the mechanism presented in this paper will be a useful step towards virtual characters who can be controlled by people while staying "in character."

Video Figure

We have included with this paper a short video showing the installation as it was presented at SIGGRAPH 2001. The video is available online at the following URL:

http://www.media.mit.edu/~badger/alphaWolf/alphaWolf.mov

Acknowledgements

The *AlphaWolf* installation has been built with the code base developed by the Synthetic Characters Group over the last five years. Therefore, we'd like to thank all the other members of the group who have contributed to our system. In addition, special thanks to Adolph Wong, who did all the animations for the virtual wolves, to Michael Patrick Johnson for his discussion of computer games, and to Roz Picard, Cynthia Breazeal and Richard Wrangham for their help with all aspects of *AlphaWolf*.

References

ALLISON, D., WILLS, B., HODGES, L., AND WINEMAN, J. 1996. Gorillas in the Bits, Georgia Institute of Technology**:** GIT-GVU-96-16.

ASSANIE, M. 2002. *Directable Synthetic Characters*. AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment.

BADLER, N., AND ALLBECK, J. 2001. Towards behavioral consistency in animated agents. *Deformable Avatars*. N. Magnenat-Thalmann, and Thalmann, D. Dordrecht, Kluwer Academic Publishers.

BATES, J., LOYALL, A., AND REILLY, W. 1992. *An Architecture for Action, Emotion, and Social Behavior*. Proceedings of the Fourth European Workshop on Modeling Autonomous Agents in a Multi-Agent World.

BLUMBERG, B., AND GALYEAN, T. 1995. *Multi-Level Direction of Autonomous Creatures for Real-Time Virtual Environments*. Proceedings of ACM SIGGRAPH 95.

BLUMBERG, B. M. 1996. Old Tricks, New Dogs: Ethology and Interactive Creatures. *Media Laboratory*. Cambridge, MA, MIT.

---

[5] P-values in this section were determined using the Binomial test.

BRAND, M. 1999. *Voice Puppetry*. Proceedings of ACM SIGGRAPH 99.

BREAZEAL, C. 2000. Sociable Machines: Expressive Social Exchange Between Robot and Human. *Artificial Intelligence Laboratory*. Cambridge, MA, MIT.

BURKE, R., ISLA, D., DOWNIE, M., IVANOV, Y. AND BLUMBERG, B. 2001. CreatureSmarts: The Art and Architecture of a Virtual Brain. *Proceedings of the Game Developers Conference*: 147-166.

CASSELL, J., BICKMORE, T., BILLINGHURST, M., CAMPBELL, L., CHANG, K., VILHJÁLMSSON, H., AND YAN, H. 1999. *Embodiment in Conversational Interfaces: Rea*. Proceedings of the CHI'99 Conference.

CASSELL, J., PELACHAUD, C., BADLER, N., STEEDMAN, M., ACHORN, B., BECKET, T., DOUVILLE, B., PREVOST, S., AND STONE, M. 1994. *Animated Conversation: Rule-Based Generation of Facial Expression, Gesture and Spoken Intonation for Multiple Conversational Agents*. Proceedings of ACM SIGGRAPH 94.

CHI, D., COSTA, M., ZHAO, L., AND BADLER, N. 2000. *The EMOTE Model for Effort and Shape*. Proceedings of ACM SIGGRAPH 2000.

DAMASIO, A. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York, G. P. Putnam's Sons.

DARWIN, C. 1965 (originally published 1872). *The Expression of the Emotions in Man and Animals*. Chicago, The University of Chicago Press.

EKMAN, P. 1992. An Argument for Basic Emotions. *Basic Emotions*. N. Stein, and Oatley, K. Hove, UK, Lawrence Erlbaum**:** 169-200.

FOX, M. W. 1971. *Behaviour of Wolves, Dogs and Related Canids*. New York, Harper & Row.

GADANHO, S., AND HALLAM, J. 1998. *Emotion-triggered Learning for Autonomous Robots*. Simulation of Adaptive Behavior 1998 Workshop on Grounding Emotions in Adaptive Systems.

HAYES-ROTH, B., SINCOFF, E., BROWNSTON, L., HUARD, R., AND LENT, B. 1995. *Directed Improvisation with Animated Puppets*. Proceedings of CHI '95.

HODGINS, J. K., AND POLLARD, N. S. 1997. *Adapting Simulated Behaviors For New Characters*. Proceedings of ACM SIGGRAPH 97.

JOHNSON, M. B. 1994. WavesWorld: A Testbed for Three Dimensional Semi-Autonomous Animated Characters. *Media Arts & Sciences*. Cambridge, MA, MIT.

JOHNSON, M. P., WILSON, A., BLUMBERG, B., KLINE, C., AND BOBICK, A. 1999. *Sympathetic interfaces: using a plush toy to direct synthetic characters*. Proceedings of CHI 99.

MAGNENAT-THALMANN, N., KALRA, P., AND ESCHER, M. 1998. *Communicating with virtual characters*. WSCG '98.

MECH, L. D., ADAMS, L. G., MEIER, T. J., BURCH, J. W., AND DALE, B. W. 1998. *The Wolves of Denali*. Minneapolis, University of Minnesota Press.

MURIE, A. 1944. The Wolves of Mount McKinley. *Fauna of the National Parks Series*. Washington, DC, U.S. National Park Service. **5**.

PERLIN, K., AND GOLDBERG, A. 1996. *Improv: A System for Scripting Interactive Actors in Virtual Worlds*. Proceedings of ACM SIGGRAPH 96.

PICARD, R. 1997. *Affective Computing*. Cambridge, MA, MIT Press.

PLUTCHIK, R. 1991. *The Emotions*. Lanham, MD, University Press of America.

REILLY, W. S. N. 1996. Believable Social and Emotional Agents. *School of Computer Science*. Pittsburgh, PA, Carnegie Mellon University.

REYNOLDS, C. 1987. *Flocks, Herds and Schools: A Distributed Behavioral Model*. Proceedings of ACM SIGGRAPH 87.

ROSE, C. 1999. Verbs and Adverbs: Multidimensional Motion Interpolation Using Radial Basis Functions. *Department of Computer Science*. Princeton, NJ, Princeton University.

ROSENFIELD, M. P., B. 1988. White Wolf. *National Geographic Video*. J. Brandenburg.

RUSSELL, J. 1997. Reading emotions from and into faces: resurrecting a dimensional-contextual perspective. *The Psychology of Facial Expression*. J. Russell, and Fernandez-Dols, J. Cambridge, Cambridge University Press.

SCERRI, P., PYNADATH, D., AND TAMBE, M. 2001. *Adjustable autonomy in real-world multi-agent environments*. Proceedings of the 5th International Conference on Autonomous Agents.

SCHENKEL, R. 1967. Submission: its features and function in the wolf and dog. *Amer. Zool.* **7**: 319-329.

SCHLOSBERG, H. 1954. Three dimensions of emotions. *Psychological Review* **61**(2): 81-88.

SMITH, C. 1989. Dimensions of appraisal and physiological response in emotion. *Journal of Personality and Social Psychology* **56**: 339-353.

STRASSMAN, S. 1994. *Semi-autonomous animated actors*. Proceedings of the 12th National Conference on Artificial Intelligence.

TERZOPOULOS, D., TU, X., AND GRZESZCZUK, R. 1994. *Artificial Fishes: Autonomous Locomotion, Perception Behaviour and Learning in a Simulated Physical World*. Artificial Life IV.

THALMANN, D., NOSER, H., AND HUANG, Z. 1997. Autonomous Virtual Actors Based on Virtual Sensors. *Creating Personalities*. R. Trappl, and Petta, P. Heidelberg, Springer Verlag.

TOMLINSON, B., AND BLUMBERG, B. 2002 (to appear). *Synthetic Social Relationships in Animated Virtual Characters*. From Animals to Animats 7: Proceedings of the Seventh International Conference on Simulation of Adaptive Behavior (SAB '02).

VELASQUEZ, J. 1998. When Robots Weep: Emotional Memories and Decision-Making. *Proceedings of the Fifteenth National Conference on Artificial Intelligence*.

Leashing the AlphaWolves:

Mixing User Direction with Autonomous Emotion in a Pack of Semi-Autonomous Virtual Characters

Bill Tomlinson, Marc Downie, Matt Berlin, Jesse Gray, Derek Lyons, Jennie Cochran, and Bruce Blumberg