AlphaWolf: Social Learning, Emotion and Development in Autonomous Virtual Agents

Bill Tomlinson and Bruce Blumberg

Synthetic Characters Group, The Media Lab, MIT
77 Massachusetts Avenue, NE18-5FL, Cambridge, MA 02139 USA
{badger, bruce}@media.mit.edu

Abstract. We present research in synthetic social behavior for interactive virtual characters. We describe a model from the natural world, the gray wolf (*Canis lupus*), and the social behavior exhibited by packs of wolves, to use as the target for an interactive installation entitled *AlphaWolf*, which was shown at SIGGRAPH 2001. We offer a computational model that captures a subset of the social behavior of wild wolves, involving models of learning, emotion and development. There is a range of real-world applications of synthetic social behavior, from short-term possibilities such as autonomous characters for computer games, to long-term applications such as computer interfaces that can interact more appropriately with humans by utilizing human social abilities. Our research offers initial steps toward computational systems with social behavior, in hope of making interactions with them more functional and more inherently rewarding.

1 Introduction

This paper presents research in social learning currently under way in the Synthetic Characters Group at the MIT Media Lab, headed by Professor Bruce Blumberg. The goal of our group is to understand the nature of intelligent behavior by creating computational entities that are inspired by animals. Over the last several years, we have developed a synthetic-character-building toolkit with which we create our virtual entities. By drawing lessons from nature, we hope to shed some light on certain hard problems in artificial intelligence, including action selection, machine learning, motor control and multi-agent

Fig. 1. A wolf pup and his father howl together

coordination. The focus of one of our current projects, entitled *AlphaWolf*, is social learning in a virtual 3D-animated wolf pack (see Fig. 1). *AlphaWolf* is an interactive installation in which people play the role of wolf pups in a pack of autonomous and semi-autonomous virtual wolves. By howling, growling, whining or barking into a microphone, each participant is able to influence how his or her pup interacts with the other members of the pack. The pups, in turn, form emotional memories that color how they will interact with their packmates in future encounters. *AlphaWolf* premiered in the Emerging Technologies section of SIGGRAPH 2001.

Various other researchers have studied synthetic social systems from natural models, including flocking in birds [29], schooling in fish [37], virtual gorillas [1], chimpanzees [35] and primate-like artificial agents [19]. Our research differs from these efforts in our focus on learning as a key component of social competence. Social learning allows creatures to benefit from context-preservation [9] – in each successive inter-character interaction, individuals can bring their interaction history to bear on their decision-making process. In the process of building this installation, we have augmented our character-building toolkit (described elsewhere – [3], [6], [13], [20]) with a variety of computational representations that are necessary to enable social learning in our virtual wolves. By taking an animal model and trying to replicate the social phenomena found in that species, we seek to create computational mechanisms that will enable social competence, learning and development in a wider range of socially intelligent agents and other computational systems.

We are pursuing two main areas of research that pertain to social learning. The first is the relationship between social learning and emotion. Our wolves feature a model of emotion and a dynamic range of behavior that is influenced by their emotional state. Wolves' emotions also play into their formation of context-specific emotional memories, which affect how they will interact in the future. Our research group has completed simple implementations of these representations, featured in the virtual wolves we showed at SIGGRAPH 2001.

The second area of research involves development as it relates to social learning. Animals are born with certain innate behaviors; suckling, for example, is too crucial a behavior to be left to learning. Simulating the timing of innate behaviors (i.e., when they become active or inactive) in the context of our learning system will allow virtual creatures to be created who are more similar to real animals. A blend of hard coded behaviors and learned behaviors will create virtual creatures with predispositions to behave in certain fashions, but the ability to adapt to the vagaries of their local environment. This is the main area of future work for our project, which we hope to implement in the coming months. Each of our wolves will play different roles as it grows up, depending on the composition of its pack, its built-in developmentally-timed behavioral predispositions and the learning it has undergone during its life.

There are a variety of real-world applications of our research. Any population of virtual or physical agents, from crowd scenes in animated movies to autonomous robots collecting samples on Mars, could benefit from computational models of social learning. Just as individual animals provide good models for creating individual agents, populations of animals may inspire mechanisms by which those individual agents will interact. Because the environments inhabited by many populations of artificial entities are variable and uncertain, it is important to have creatures who are predisposed toward certain behaviors but can learn and adapt to meet the constraints imposed by their environment.

2 The Gray Wolf

In their natural environment, gray wolves form complex social groups called 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. [26], [23] The average pack size is approximately 7-9 individuals, but some packs may contain more than 25 wolves. Large packs may contain more than one breeding pair. Most young wolves disperse from their natal pack in their second or third year to find a mate and begin their own pack. [23]

Wolves communicate with each other in a variety of ways. They have a wide array of vocalizations, including "whimpering, wuffing, snarling, squealing and howling". [40, p. 68] They express their intentions and motivational and emotional states through body posture as well – a mother wolf assumes different postures with her pups than she does with her mate. The sense of smell is also integral to wolf social behavior (e.g., scent marking). In wolves, as in most social creatures, communication is central to the social relationships that are formed.

When the pack has young pups, the adult wolves travel away from the den to hunt, and carry back meat in their stomachs to feed the pups. Upon their return, the pups perform stereotypical food-begging behavior, in which they crouch in front of an adult and lick or peck at the adult's muzzle. This pup behavior incites the adult to regurgitate the meat, which the pups excitedly consume. [32]

Wolf social behaviors appear to be derived from other behavioral patterns exhibited by wolves. [32] For example, 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. These behavioral patterns bear a resemblance to infantile behaviors involved in reflex urination (in which a pup urinates when his mother licks his belly). [14] Active submission involves a crouched posture with backward directed ears, and licking or pecking the mouth of the dominant wolf. This behavior is very similar to the food-begging behavior of pups described above. Similarly, dominant behaviors appear to be a form of "ritualized fighting". [17]

We chose the gray wolf as the model for our simulation for several reasons. First, they manifest distinct social phenomena that are complex enough to be interesting, yet clear enough to provide direction for our simulation. Second, wolves are closely related to the domestic dog, for which we have a strong conceptual and technical base as a result of previous installations that we have done featuring virtual dogs. Finally, the social behaviors of wolves are similar enough to those of humans that some of the lessons we learn from wolves might be relevant to human social behavior and simulation.

3 Computational Representations: Social Learning and Emotion

There are a variety of computational elements that must be in place for virtual wolves to interact socially in a way resembling real wolves. Our virtual wolves must be able to choose different behaviors; to move around their world; to learn that certain interactions lead to positive results while others lead to negative repercussions. These components are already functional parts of our character-building toolkit, and have been described elsewhere. [20], [13], [6], [3]

The fundamental representation of an action in our system is the ActionTuple. Each ActionTuple consists of four components: the action itself; a TriggerContext, which determines when the action will take place; a DoUntilContext that determines when the action will cease; and an object to which the action will happen. ActionTuples are arranged into ActionGroups that determine which actions are mutually exclusive and which can be run simultaneously. In accord with Thorndike's Law of Effect [36], ActionTuples that are correlated with positive results will be chosen more frequently in the future. Each action is executed in an emotional style, which we discuss below.

Social learning is intimately linked with the ability to have emotions, to express those emotional states, and to remember an association between environmental stimuli and emotional states. The main extensions to our system that have already been implemented for this project are: the ability to have and express emotional states, and the ability to form context-specific emotional memories. We will describe each of these elements in greater depth below.

3.1 Having Emotions

In order to simulate emotional virtual characters, it is necessary to choose a computational representation that captures the range of emotional phenomena we wish our characters to exhibit. Much research has already been done both in understanding emotions and in simulating them computationally. Darwin's ideas about emotions [11] form the basis for much of modern research into understanding emotions scientifically.

For the *AlphaWolf* project, we have considered two main emotional paradigms – a dimensional approach and a categorical approach. The dimensional approach (e.g., [33], [28], [31], [34]) maps a range of emotional phenomena onto explicitly dimensioned space. Various researchers have implemented versions of the dimensional approach; for example, Breazeal [4] maps a 3-dimensional space (Arousal, Valence, Stance) onto a set of emotions that influence both the behavior system and the motor system.

The categorical approach separates emotional phenomena into a set of basic emotions – for example, fear, anger, sadness, happiness, disgust and surprise. [14] Ekman's model provided the basis for an implementation by Velasquez [38]; others (e.g., [16]) have also implemented categorical models. For a far more comprehensive discussion of emotional models in computational systems, the reader is directed to Rosalind Picard's book, *Affective Computing* [27].

A dimensional approach better captured the range of emotional phenomena that we wanted our wolves to display. Our emotion model is based most directly on the Pleasure-Arousal-Dominance model presented by Mehrabian and Russell [24]. At each moment, a wolf has three continuous variables describing a 3-dimensional emotional space. Emotional categories may also be mapped onto these axes; for example, anger is equivalent to low Pleasure, high Arousal and high Dominance. The value for each variable is affected by the wolf's previous emotional state, by some drift that each emotional axis undergoes, and by attributes of his surrounding environment. Currently our Dominance axis is most thoroughly fleshed out; we hope to elaborate on the other two axes in the coming months.

Fig. 2. The wolf pup play-bows at his father

3.2 Expressing Emotional States

The emotional state of our virtual wolves affect the action selection mechanism by feeding into the TriggerContexts. For example, a wolf's dominance value (D, where $0 < D < 1$) will factor into a trigger that causes growling to occur, and $(1 - D)$ will factor into the trigger that makes the wolf submit.

In addition, the emotional states of our wolves feed into our motor control system [13] and affect the style in which they take their actions. This system is based on the "verbs and adverbs" system of Rose [30], in which an action (a "verb") is taken in a certain style (an "adverb"). Because of our motor control system, which can blend between example animations, our animators need only create a few extreme emotional styles of each action (for example, one high Arousal walk and one low Arousal walk) to get the full dynamic expressive range between those examples. Finally, emotion is central to our creatures' social interactions.

3.3 Context-Specific Emotional Memories

A component of real wolf social behavior is that submissive wolves know to submit *before* the dominant wolf even reaches them, and thus avoid a potentially harmful interaction. [32] Similarly, our virtual wolves need some way of triggering submissive behaviors before they are actually pinned to the ground by a dominant individual. How can our wolves remember previous interactions that they have had? In our simulation, we utilize *context-specific emotional memories* (CSEMs) for this purpose. CSEMs cause a wolf who is experiencing a suite of environmental stimuli (for example, a dominant wolf) to return to an emotional state similar to the one he experienced on previous encounters with those stimuli. In addition to storing the values of Pleasure, Arousal and Dominance that the wolf was feeling when he last experienced the relevant stimuli, the CSEM also features a Confidence value that reflects how reliable he believes the CSEM's values to be. Each continuously changing CSEM effectively reflects the interaction history between the wolf and some bit of his context, without the need for specific memories of past interactions.

Our CSEMs are based on the "somatic marker hypothesis" presented by Damasio [10], in which he proposes that people attach emotional significance to stimuli that they encounter in their environment, and then re-experience that emotion when they encounter those stimuli on future occasions. Other researchers have implemented models of emotional learning or memory, for example "affective tags" [39], "emotional memories" [38], and others (e.g., [21], [16], [2]).

An important extension to our system would be to allow our wolves to form multiple CSEMs for one object, depending on the context in which the wolf has encountered that object. For example, a wolf pup might form one CSEM about his father when the father smells of meat (i.e., "He'll probably regurgitate food for me if I harass him!"), and a very different one when he doesn't smell of meat (i.e., "He'll probably pin me to the ground if I harass him!") Both of these CSEMs would have high values for Arousal, but very different values for Pleasure.

4 Future Work: Social Learning and Development

A primary area of research that we hope to pursue in the near future involves giving our wolves the ability to learn to use behaviors in novel social contexts. We mentioned above that wolf submission behaviors are similar to infantile behaviors. How do wolves learn to use these behaviors to negotiate their social relationships? Over the next few paragraphs, we will describe some of the elements of our learning system that pertain to this co-opting of one behavior for another purpose.

In our previous project, entitled "Clicker by Eire", our virtual terrier Duncan had the ability to learn to exhibit a behavior in response to an appropriate context. [3] A human participant could train Duncan to do a variety of tricks in response to voice commands. Duncan's learning system back-propagated the value of a food treat to reinforce actions that he had been taking during a certain time window that preceded the treat. In addition, Duncan was able to distinguish between contexts in which he was rewarded and those in which he was not rewarded, and only propagated the value of the treat to those contexts in which a reward occurred. This is the essence of the mechanism by which our wolves will learn to perform certain behaviors in the correct context.

Here is an example of how this learning might proceed in our wolves. Imagine a pup and a father wolf who habitually eat from the same carcass. As the pup gets older, the father will become less tolerant of the pup grabbing food away from him. At a certain point, the father may discipline the pup for eating too close to him, by growling and dominating the pup. The act of disciplining will continue until the pup submits or escapes. If submission behavior tends to cause the father wolf to relent, the pup should rapidly learn to submit in the context of being disciplined. Fleeing will also separate the pup from the noxious stimulus that is an annoyed father wolf.

By this mechanism, the pup will learn to submit or flee in the presence of a growling adult.

A related behavioral component of our virtual wolves is that dominant individuals need to be willing to relent once the submissive individual actually submits. Possibilities include a hard-coded "evolutionary taboo" against hurting things that act like pups or a learned association that a submitting individual will no longer perform whatever action caused the dominant's distress.

The *AlphaWolf* project is changing the way we build characters – rather than assembling finished adults, we are creating young pups with certain built-in behaviors and allowing them to grow up in a social context where they can re-use and modify these behaviors for other purposes. Additional suites of behaviors may be scheduled to "come online" at certain points in development. For example, at sexual maturity, a whole new group of behaviors become eligible for expression. In addition to changing behaviorally as they age, the wolves also change anatomically, their physical forms growing and morphing from pup to adult.

Change on developmental time scale is finding its way into both academic research and commercial products. For example, the computer game "Creatures" [18] and other virtual pets grow up as you interact with them. Cañamero's Abbotts project [7] also features a developmental perspective.

5 Applications

While our group's research is not explicitly directed toward any specific application, we often discuss our work with people who have some interest in commercial, industrial or academic applications. Through these discussions, we have considered a variety of applications for socially competent computational systems. While our group does not work on implementing these ideas, we present them here as potential reasons why socially intelligent agents might be useful research.

The entertainment industry might have a use for socially competent agents. Computer games would be more fun to play if the "good guys" and "bad guys" were able to form relationships with each other and with the players. Animated agents could help make movies by interacting socially to create more realistic crowd scenes and background action. Toys that could learn about the kids who play with them might be more engaging to those kids.

Socially competent agents could serve as an instructional tool for teaching social skills to a variety of age groups, from children to corporate managers. For example, the *AlphaWolf* installation makes social relationships explicit. When a participant growls, his pup growls. When he growls *at* another wolf, that wolf becomes more submissive. In addition, the wolf he growled at will be more likely to submit to him in the future. By making a clear connection between social action and social effect, the *AlphaWolf* project or some other experience involving social synthetic characters could make it easier for people at all stages to become aware of their own social behavior.

Any complex ecology of entities must have some mechanism for introducing new members, allocating resources, and resolving conflicts. Wild wolf populations do this, with new pups integrating into the pack each season, dominance relationships aiding in resource allocation, and status signals helping to prevent conflicts by mutual consent. Perhaps there are lessons to be learned from natural social organizations of heterogeneous (e.g., different ages) entities, which can be applied to systems created by humans.

Finally, since humans are such social creatures, there might be some use for computational models of social behavior in human-computer interfaces. Building interfaces that are able to take advantage of humans' natural social abilities will make the systems behind those interfaces both more functional and more inherently rewarding.

6 Conclusion

Social behavior is often evolutionarily advantageous in the natural world. Research in synthetic social agents allows us to harness the benefits of natural social behavior for use in computational systems. Taking one species, in this case the gray wolf, and using it as a model for our simulation allows us to have a yardstick by which to measure our success.

The essential mechanism by which our wolves form social relationships is this: each wolf has an emotional state at every moment; the wolf is able to recognize all of the other wolves; he is able to form an association between each other wolf and the emotional state that he tended to experience during previous interactions; finally, that emotional association influences his current emotional state when he again encounters a wolf for whom he has formed an emotional memory. This simple mechanism allows our wolves to remember each other and pick up where they left off in their social interactions.

Potential real world applications of synthetic social behavior are as diverse as the applications of real social entities. From characters in computer games, to educational aids for social skill development, to mechanisms for conflict negotiation in virtual or robotic entities, social awareness and interactions may play a significant role in a range of industries.

We have just exhibited the *AlphaWolf* installation at SIGGRAPH 2001, where more than a thousand people interacted with our virtual wolves. While we did not collect user data at SIGGRAPH, most participants appeared to enjoy playing with socially competent virtual wolves. People frequently brought their friends back to see the wolves, describing to them the behavior that their pups exhibited. Participants were very willing to become engaged with our wolf pups, immersing themselves in the interactive experience.

Humans are social animals; our computational systems should be able to engage our social abilities. The research described in this paper helps address how we can make computational entities who can learn to have social relationships with each other. This is a small step toward making systems that can take advantage of the natural benefits of social organization, and an even smaller step toward introducing computers to the complex world of human social interactions.

7 Acknowledgements

We would like to thank all the members and friends of the Synthetic Characters Group who have helped make our wolves possible: Marc Downie, Matt Berlin, Jesse Gray, Adolph Wong, Robert Burke, Scott Eaton, Damian Isla, Yuri Ivanov, Michael Patrick Johnson, Spencer Lynn, Ben Resner, Derek Lyons, Jennie Cochran, Bryan Yong, Steve Curcuru, Geoffrey Beatty, Jed Wahl, Dan Stiehl, Rusmin Soetjipto, Dan Zaharopol, Aileen Kawabe, Madaline Tomlinson and Professor Irene Pepperberg. In addition, many of the ideas in this paper were fleshed out through discussions with Professors Rosalind Picard and Richard Wrangham, and with the participants in SIGGRAPH 2001 who interacted with the *AlphaWolf* installation and with us.

References

1. Allison, D., Wills, B., Hodges, L., and Wineman, J.: Gorillas in the Bits, Technical Report, GIT-GVU-96-16, Georgia Institute of Technology. (1996)
2. Balkenius, C. and Moren, J.: A Computational Model of Context Processing, in SAB '00, (2000) 256-265
3. Blumberg, B.: Building characters that learn what they ought to learn. International Workshop on Story Telling. Avignon, France.(2001)
4. Breazeal, C.: Sociable Machines: Expressive Social Exchange Between Robot and Human. Ph.D. Thesis, Artificial Intelligence Laboratory, MIT.(2000)
5. Buck, R.: The Communication of Emotion. New York, NY: Guilford Press. (1984)
6. Burke, R., Isla, D., Downie, M., Ivanov, Y. and Blumberg, B.: CreatureSmarts: The Art and Architecture of a Virtual Brain, in Proceedings of the Game Developers Conference, (2001) 147-166
7. Cañamero, D.: Modeling Motivations and Emotions as a Basis for Intelligent Behavior. In Proceedings of 1st Int'l. Conf. on Autonomous Agents (1997)
8. Cañamero, D.: Issues in the Design of Emotional Agents. In Emotional and Intelligent: The Tangled Knot of Cognition. Papers from the 1998 AAAI Fall Symposium. Menlo Park, CA: AAAI Press.(1998) 49-54
9. Cohen, M.D., Riolo, R.L., and Axelrod, R.: The Emergence of Social Organization in the Prisoner's Dilemma: How Context-Preservation and Other Factors Promote Cooperation. Santa Fe Institute Working Paper 99-01-002. (1999)
10. Damasio, A.: Descartes' Error: Emotion, Reason, and the Human Brain. New York, NY: G.P.Putnam's Sons. (1994)
11. Darwin, C.: *The Expression of the Emotions in Man and Animals*. Chicago: The University of Chicago Press. Originally published in 1872. (1965)
12. Dautenhahn, K.: Socially Intelligent Agents and The Primate Social Brain – Towards a Science of Social Minds. In Socially Intelligent Agents: The Human in the Loop. Papers from the 2000 AAAI Fall Symposium. Menlo Park, CA: AAAI Press. (2000) 35-51
13. Downie, M.: Behavior, Animation and Music: The Music and Movement of Synthetic Characters. M.S. Thesis, Media Lab, MIT. (2001)
14. Ekman, P.: An Argument for Basic Emotions. Stein, N. and Oatley, K. eds. Basic Emotions. Hove, UK: Lawrence Erlbaum. (1992) 169-200
15. Fox., M.: The Wild Canids. Van Nostrand Reinold. (1975)
16. Gadhano, S. and Hallam, J.: Emotion-triggered Learning for Autonomous Robots. in SAB '98 Workshop on Grounding Emotions in Adaptive Systems. (1998)

17. Golani, I. and Moran, G. A motility-immobility gradient in the behavior of the "inferior" wolf during "ritualized fighting". in Eisenberg, J. F. and D. G. Kleiman, ed. *Advances in the study of mammalian behavior*. The American Society of Mammalogists, Special Publication No.7. (1983) 65-94
18. Grand, S., Cliff, D., and Malhotra, A. Creatures: Artificial Life Autonomous Software Agents for Home Entertainment. in In Proceedings of 1st Int'l. Conf. on Autonomous Agents (1997).
19. Hemelrijk, C.: An individual-orientated model of the emergence of despotic and egalitarian societies. Proc. R. Soc. Lond. B 266 (1999) 361-369
20. Isla, D., Burke, R., Downie, M. and Blumberg, B.: A Layered Brain Architecture for Synthetic Characters. IJCAI (2001)
21. Kitano, H.: A Model for Hormonal Modulation of Learning. in IJCAI (1995) 532-538
22. Mech, L.: Alpha status, dominance, and division of labor in wolf packs. Canadian Journal of Zoology 77 (1999) 1196-1203
23. Mech, L D., Adams, L. G., Meier, T. J., Burch, J. W., and Dale, B. W.: The Wolves of Denali. Minneapolis, MN: University of Minnesota Press. (1998)
24. Mehrabian, A. and Russell, J.: An Approach to Environmental Psychology. Cambridge, MA: MIT Press.(1974)
25. Morss, C.: Seminar on wolf behavior. Wolf Hollow Ipswich, Ipswich, MA. (Fall 2000)
26. Murie, A.: The Wolves of Mount McKinley. Fauna of the National Parks Series, No. 5. Washington, DC.: U.S. National Park Service. (1944)
27. Picard, R.: Affective Computing. Cambridge, MA: MIT Press. (1998)
28. Plutchik, R.: The Emotions. Lanham, MD: University Press of America. (1991)
29. Reynolds, C.: Flocks, Herds and Schools: A Distributed Behavioral Model. Computer Graphics, 21(4). In Proceedings of the SIGGRAPH '87 Conference (1987) 25 - 34
30. Rose, C.: Verbs and Adverbs: Multidimensional Motion Interpolation Using Radial Basis Functions. Ph.D. Dissertation, Department of Computer Science, Princeton University. (1999)
31. Russell, J.: Reading emotions from and into faces: resurrecting a dimensional-contextual perspective. in Russell, J. and Fernandez-Dols, J. (eds.) The Psychology of Facial Expression. Cambridge: Cambridge University Press. (1997)
32. Schenkel, R.: Submission: its features and function in the wolf and dog. American. Zoologist. 7 (1967) 319-329
33. Schlosberg, H.: Three dimensions of emotions. Psychological Review, 61(2) (1954) 81-88
34. Smith, C.: Dimensions of appraisal and physiological response in emotion. Journal of personality and social psychology, 56 (1989) 339-353
35. teBoekhorst, I. and Hogeweg, P.: Self-structuring in artificial 'CHIMPS' offers new hypothesis for male grouping in chimpanzees. Behaviour 130 1994) 229-52
36. Thorndike, E.: Animal Intelligence. Darien, CT: Hafner. (1911)
37. Tu, X. and Terzopoulos, D.: Artificial Fishes: Physics, Locomotion, Perception, Behavior. In Proceedings of ACM Computer Graphics, SIGGRAPH'94. (1994) 43-50
38. Velasquez, J.: When Robots Weep: Emotional Memories and Decision-Making. In Proceedings of the Fifteenth National Conference on Artificial Intelligence. Madison, Wisconsin: AAAI Press. (1998)
39. Yoon, S., Blumberg, B. and Schneider, G.: Motivation Driven Learning for Interactive Synthetic Characters. In Proceedings of 4th Int'l. Conf. on Autonomous Agents (2000)
40. Zimen, E.: The Wolf. New York, NY: Delacorte Press. (1981)