Social Behavior, Emotion and Learning in a Pack of Virtual Wolves

Bill Tomlinson and Bruce Blumberg

Synthetic Characters Group, The Media Lab, MIT
E15-320G, 20 Ames Street, Cambridge, MA 02139 USA
{badger | bruce}@media.mit.edu

Abstract

We are creating a pack of virtual creatures who exhibit the kinds of social interactions found in a natural species of animal, the gray wolf (*Canis lupus*). To do this, we are extending our synthetic character building toolkit to enable our characters to learn to adapt pre-existing behaviors for use in novel social contexts; to have and express emotional states; and to form context-specific emotional memories. We describe how these elements combine to form the underpinnings for our interactive installation entitled "AlphaWolf", to be shown at SIGGRAPH 2001. We believe that the computational representations that allow social learning in our virtual wolves demonstrate the intimate connections among social behavior, emotion and learning. In addition, we feel our findings are applicable to building a wider range of socially intelligent agents. In addition, we hope that our studies of virtual wolves will offer insight into the processes by which real wolves and other animals understand their environments.

Introduction

"The submissive activity is, in its essence, an activity of the cub."

Rudolf Schenkel, Submission: its features and function in the wolf and dog. *American Zoologist.* 1967, p. 325

Social behavior is often evolutionarily advantageous in the natural world. Many species of animals live part or all of their lives in social groups, from the transient annual assemblages of elephant seals at their breeding grounds to the cities and towns of human cultures. Various researchers have studied synthetic social systems, including flocking in birds (Reynolds 1987), schooling in fish (Tu & Terzopoulos 1994), virtual gorillas (Allison et al. 1996), chimpanzees (teBoekhorst & Hogeweg 1994) and primate-like artificial agents (Hemelrijk 1999). This research is interesting not only because it helps us understand how nature works, but also because it may help us harness the benefits of social behavior for use in computational systems.

One species of animal that is renowned for its social organizations is the gray wolf (*Canis lupus*). Wolves exhibit an array of interesting social phenomena – they live and hunt in packs, they display hierarchical social relationships that negotiate the allocation of scarce resources, they communicate through a variety of sensory channels (e.g., howling, scent marking, body posture). In the Synthetic Characters Group at the MIT Media Lab, we are in the process of creating a pack of virtual wolves (see Figure 1) who exhibit similar social behavior to real wolves.

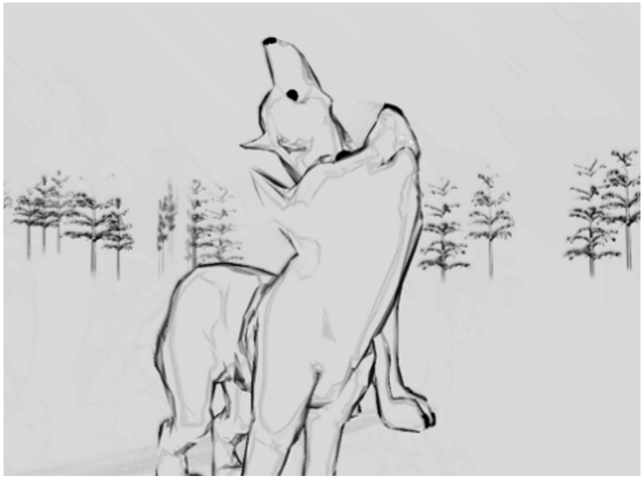

Figure 1: A wolf pup and his father howl together.

We are giving our wolves a variety of social competences, including the ability to learn to use behaviors in novel social contexts, to express their emotional states and to respond to other wolves in socially appropriate ways. These social abilities are made possible by our models of action-selection, emotion, motor control, learning and context-specific emotional memory formation. By taking an animal model and trying to replicate the social phenomena found in that species, we hope to shed light on computational mechanisms that could enable social competence, learning and development in a wider range of socially intelligent agents (SIAs, from Dautenhahn 2000) and other computational systems. In addition, our experiments in simulation may be of some use to the biological community in framing their experiments with real animals.

In order to ensure that our work accurately represents the biology and ethology of wild wolves, we begin this paper with a description of wolf biology and social behavior. Next, we analyze the computational representations that we believe need to be in place to enable social learning in our virtual wolves, with consideration of the biological literature and other models, both theoretical and computational, that have been proposed. Following that, we discuss an interactive installation that is currently under development as a result of this research project. We close with some ideas for future work and a consideration of how this work can impact the broader field of socially competent computational systems.

The work described in this paper is made possible by a range of other research that is taking place simultaneously in the Synthetic Characters Group. While there is not space to describe each project here, we nevertheless want to mention research in behavior system design (Isla et al. 2001), motor control (Downie 2001), and learning (Blumberg et al. forthcoming), all of which are essential to our wolf social simulation.

Animals are the world's best example of autonomous creatures; animals, therefore, provide the inspiration for our research in virtual autonomous creatures. 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. In addition, our simulation could be used to inform the study of the biological species that inspired it. These premises all reflect the "Life-Like Agents Hypothesis" proposed by Dautenhahn – "Artificial social agents (robotic or software) which are supposed to interact with humans are most successfully designed by imitating life." (1999, cited in Dautenhahn 2000, p.37)

The Gray Wolf

In order for readers of this paper to be able to judge our ideas about simulated wolves, we begin by presenting the elements of real wolf behavior that we find so fascinating. Much of this information is drawn from the research of David Mech. (Mech et al. 1998)

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 (Murie 1944, Mech et al. 1998). The average pack size is approximately 7-9 individuals, but some packs may contain more than 25 wolves. Larger 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. (Mech et al. 1998)

Wolves communicate with each other in a variety of ways. They have a wide array of vocalizations, including "whimpering, wuffing, snarling, squealing and howling". (Zimen 1981, 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. Scent marking is the most well known kind of olfactory communication in wolves, relating to the lifted-leg urination found in domestic dogs. In wolves, as in most social creatures, communication is central to the social relationships that are formed.

Wolves may hunt singly or in packs. Pack hunting allows wolves to take down prey that would be too big for a single wolf to catch – for example, caribou or moose. When the pups are too young to travel, the adults in the pack travel away from the den to hunt, and, after a successful hunt, carry back meat in their stomachs to feed the pups. Upon their return, their 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. (Schenkel 1967)

Wolf social behaviors appear to be derived from other behavioral patterns exhibited by wolves. (Schenkel 1967, Fox 1971) 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). 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" (Golani and Moran 1983).

We chose the gray wolf as the model for our simulation for several reasons. First, they manifest social behavior that is complex enough to be interesting, yet simple enough (and well-enough understood) for us to have a chance of capturing its essence in simulation. Second, wolves are closely related to the domestic dog, for which we have a strong conceptual and technical base as a result of our two preceding installations, which dealt with dogs. Finally, the social behaviors of wolves are not so dissimilar to those of humans that we hope some of the lessons we learn from wolves might be relevant to human social behavior and simulation.

Computational Models

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 behaviors 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. (Isla et al. 2001) (Downie 2001) (Blumberg et al. forthcoming) The main extensions to our characters that are unique to this project are: the ability to learn to use behaviors in novel social contexts; the ability to have and express emotional states; and the ability to form context-specific emotional memories. We will describe each of these in greater depth below.

Using 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 at a certain time. (Blumberg et al. forthcoming) In that installation, a human participant could train Duncan to do a variety of tricks in response to voice commands. To get him to do this, the participant needed to reward Duncan with a virtual food treat in order for him to know that he had taken the correct action. Duncan's learning system back-propagated the value of the food treat to 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 learn to propagate the value of the treat to reinforce only 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. Imagine an example involving a father wolf trying to fall asleep. If his puppy starts yapping at him (see Figure 2), the father might get up and discipline the pup. The father's disciplining might continue until he is satisfied, perhaps when the pup submits. If submission behavior tends to cause the father wolf to relent, the pup should rapidly learn to perform that kind of behavior in the context of being disciplined.

A related behavioral component of our virtual wolves is that dominant individuals need to be willing to relent once the submissive individual actually submits. Since the dominant individual provides the reward function from which the submissive learns, it is necessary for that dominant individual to have some computational mechanism that causes him to relent. Possibilities include a hard-coded "evolutionary taboo" against hurting things that act like pups or a learned behavior that a submitting individual will no longer perform whatever action caused the dominant's distress.

The ability of our virtual wolves to adapt behaviors to serve in social interactions suggests a new way of designing virtual characters – rather than assembling finished adult characters, we are creating young characters who have 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 kick in. Our software development process is becoming more like a biological development process, where the time of onset and offset of behaviors affects how those behaviors will be rolled into the behavior system.

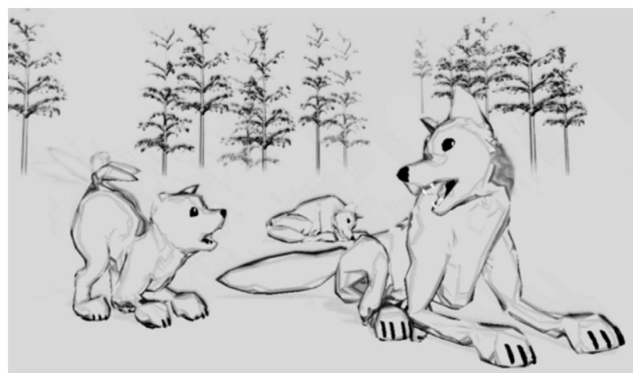

Figure 2: The wolf pup play-bows at his father.

Having and Expressing Emotional States

At least three main paradigms for conceiving of emotions have been proposed. The cognitive appraisal theory (the "OCC model") was offered by Ortony, Clore and Collins (1988) and a version of it was implemented computationally by Elliott (1994). Ekman (1992) offered a set of basic emotions – fear, anger, sadness, happiness, disgust and surprise. Ekman's model has been implemented by Velasquez (1998). A third kind of model represents emotions through an explicitly dimensioned space (e.g., (Schlosberg 1954)). Breazeal (2000) maps a 3-dimensional space (Arousal, Valence, Stance) onto a set of emotions, which in turn influence both the behavior system and the motor system. For a far more comprehensive discussion of emotional models in computational systems, the reader is directed to Rosalind Picard's book, "Affective Computing" (1997).

Our emotion model is based on the Pleasure-Arousal-Dominance model presented by Mehrabian and Russell (1974). At each moment, a wolf has three continuous values describing his emotional state, each of which represents an orthogonal axis in a 3-dimensional emotional space – one axis (Pleasure) defining his emotion as good or bad, a second (Arousal) that varies from excited to bored, and a third (Dominance) that varies from dominant to submissive. These axes map well onto the kinds of phenomena found in wolves that we want to show.

The emotional state of a wolf at a given moment is affected by his previous emotional state, by the built-in rate of drift that each element undergoes, and by attributes of his surrounding environment. Returning to the example above (see Figure 2), a drowsy father wolf's Arousal might be fairly low, and drifting lower as he falls asleep, but might be forced higher by the insistent attentions of his puppy.

The emotional states of our wolves feed into our motor control system (Downie 2001) and affect the style in which they take their actions. This system is based on the "verbs and adverbs" system of Rose (1999), in which an action (a "verb") is taken in a certain style (an "adverb"). While the actions taken by the individual are determined by its action-selection system (which may also be influenced by its emotional state), the adverbs that modify those actions are derived directly from the individual's emotional state. For example, as the Arousal of the drowsy wolf climbs as a result of the pup's harassment, his Pleasure will drop, and when he finally gets up to discipline the pup, his actions will be colored by his low Pleasure. As a result, he might be snarling and stiff-legged as he approaches the insolent pup. Because of our motor control system, our animators need only create the 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.

In her 1998 paper, Dolores Cañamero asks: "What does this particular (type of) agent need emotions for in the concrete environment it inhabits?" (p.54) The answer to this question is twofold: first, as we have just mentioned, the emotional state allows the virtual creature to express itself to other individuals in its world. This has the beneficial side effect that a person watching the creatures interact can also tell what emotional state is motivating the creature's actions. The second part of the answer will be addressed more in the next section – emotion is central to the mechanism we propose by which a creature remembers its relationship with another individual. Thus, emotion is central to our creatures' social interactions.

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. (Schenkel 1967) 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, this will involve having *context-specific emotional memories* (CSEMs) through which the presence of some environmental stimuli (for example, a dominant wolf) will cause the wolf to return to a similar emotional state to the state he was in the last time he experienced that combination of stimuli. In addition to storing the values of Pleasure, Arousal and Dominance that the wolf was feeling when he last experienced that suite of 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.

These CSEMs are based on the "somatic marker hypothesis" presented by Damasio (1994), 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. CSEMs are similar to, but distinct from, the "affective tags" presented by Yoon et al. (2000) and the "emotional memories" of Velasquez (1998). Our model differs from these implementations in one crucial way: our wolves may 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 for the father 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. Another implementation in which affective state influences the connection between memories and behavior is the "mood-congruent recall" presented by Hudlicka (1998).

In order for CSEMs to work correctly, it is important to have a model of emotion that continues to affect behavior even when the creature's emotion is at a fairly low level. A low but nevertheless active level of emotion is necessary to update the CSEMs created for all the mundane objects in the wolf's world. For example, continued exposure to trees should create a CSEM towards trees that has a very high confidence, and very low arousal (e.g., trees are usually present and rarely correlate to anything interesting). This will help the wolf focus on the important things in his world by directing his attention away from trees, rocks and other persistently uninteresting things, allowing him instead to key in on the more salient and novel elements in his perceptual fields. It will also allow him not to develop a phobia about trees just because trees happen to be nearby when something bad happens to him.

Future Work

We have three main extensions to this system that we hope to pursue as well. The first is the mechanism by which our wolves perceive the emotional states of other individuals. It would be interesting to enable wolves to visually, acoustically and olfactorily discriminate among dominant and submissive signals coming from other wolves. This has one striking conceptual repercussion: perhaps much of the learning that goes on in wolves, by which they co-opt non-social behaviors to serve as social signals, is directed by the perceptual mechanisms of other wolves. For example, if an individual is less likely to act aggressive toward a larger wolf, then learning ways of looking big (e.g., raising hackles, erecting ears and tail, standing up tall) is an excellent way to inhibit aggression from others.

The second major area of extension is in the topic of alliance formation. In captive packs, wolves will sometimes appear to work together to achieve dominance (Morss, pers. comm.) This could be a natural result of our context-specific emotional memory mechanism, if wolves form memories that relate to the simultaneous presence of multiple wolves. With respect to our learning terminology, it could be a form of state-space discovery (Blumberg et al. forthcoming) for the wolves to learn that others tend to work as teams. While this alliance formation might happen automatically as a result of the system we are building, it would need to be tested to see how it compares to the biological literature on the topic of alliance formation in wolves. This could be a very interesting extension in that it might shed some light on how alliances are formed among people, as well.

An interesting experiment for our research project would be to model wolves who exhibit the kinds of social behavior seen in the wild, and then to take several of those wolves and lock them together in close contact to see if they develop the same behavioral patterns found in captive wolves. While aggressive dominance conflicts are not uncommon in captive packs of wolves (Schenkel 1967) (Zimen 1981), they appear to be a far less significant part of wolf social life in the wild. (Mech 1999) This might suggest that dominance conflicts are to a certain extent a pathological result of the close contact enforced by captivity. We could perform experiments in our simulation to determine what elements of captivity are most responsible for the pathological behavior. We will have succeeded in our research if our experiments lead to more humane treatment of real captive wolves.

Conclusion

All of the work we are doing on wolves could have several significant intellectual results. First, it could inform the way in which we design socially and emotionally intelligent agents, helping to flesh out the ways in which agents should communicate, learn and develop with respect to each other. Second, it could be used to reflect back on wild wolves, and perhaps make suggestions to biologists about experiments that they might do, why certain misconceptions exist (such as the extreme dominance and submission exhibited in captivity but not in the wild that we mentioned earlier), and how wolves might think about their social environment. Finally, by considering the simulation of social phenomena in the simpler case of wolves, we might be able to come to a deeper understanding of the social phenomena that occur in the more complex case of humans.

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, 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 Professor Rosalind Picard and Professor Richard Wrangham.

References

Allison, D., Wills, B., Hodges, L., and Wineman, J. 1996. Gorillas in the Bits, Technical Report, GIT-GVU-96-16, Georgia Institute of Technology.

Blumberg, B., Burke, B., Downie, M., Isla, D., and Ivanov, Y. Behavioral Adaptation, State-Space Discovery and Motor Learning for Synthetic Characters. Forthcoming.

Breazeal, C. 2000. Sociable Machines: Expressive Social Exchange Between Robot and Human. Ph.D. Thesis, Artificial Intelligence Laboratory, MIT.

Buck, R. 1984. *The Communication of Emotion*. New York, NY: Guilford Press.

Cañamero, D. 1998. Issues in the Design of Emotional Agents. In *Emotional and Intelligent: The Tangled Knot of Cognition. Papers from the 1998 AAAI Fall Symposium*. 49-54. Menlo Park, CA: AAAI Press.

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

Dautenhahn, K. 2000. 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*. 35-51. Menlo Park, CA: AAAI Press.

Downie, M. 2001. Behavior, Animation and Music: The Music and Movement of Synthetic Characters. M.S. Thesis, Media Lab, MIT.

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

Elliott, C. 1994. Components of a two-way emotion communication between humans and computers using a broad, rudimentary model of affect and personality. *Cognitive Studies: Bulletin of the Japanese Cognitive Science Society*. 1(2):16-30.

Hemelrijk, C. 1999. An individual-orientated model of the emergence of despotic and egalitarian societies. *Proc. R. Soc. Lond. B* 266: 361-369.

Isla, D., Burke, R., Downie, M. and Blumberg, B. 2001. A Layered Brain Architecture for Synthetic Characters. *IJCAI 2001.*

Hudlicka, E. 1998. Modeling Emotion within Symbolic Cognitive Architecture. In *Emotional and Intelligent: The Tangled Knot of Cognition. Papers from the 1998 AAAI Fall Symposium*. 49-54. Menlo Park, CA: AAAI Press.

Mech, L. 1999. Alpha status, dominance, and division of labor in wolf packs. *Canadian Journal of Zoology* 77:1196-1203. Jamestown, ND: Northern Prairie Wildlife Research Center Home Page. (Version 16MAY2000) http://www.npwrc.usgs.gov/resource/2000/alstat/alstat.htm

Mech, L D., Adams, L. G., Meier, T. J., Burch, J. W., and Dale, B. W. 1998. *The Wolves of Denali*. Minneapolis, MN: University of Minnesota Press.

Mehrabian, A. and Russell, J. 1974. *An Approach to Environmental Psychology*. Cambridge, MA: MIT Press.

Morss, C. 2000. Seminar on wolf behavior. *Wolf Hollow Ipswich*, Fall 2000. Ipswich, MA.

Murie, A. 1944. *The Wolves of Mount McKinley*. Fauna of the National Parks Series, No. 5. Washington, DC.: U.S. National Park Service.

Ortony A., Clore, G. and Collins, A. 1988. *The Cognitive Structure of Emotions*. Cambridge, UK: Cambridge University Press.

Picard, R. 1998. *Affective Computing*. Cambridge, MA: MIT Press.

Reynolds, C. 1987. Flocks, Herds and Schools: A Distributed Behavioral Model. Computer Graphics, 21(4). In *Proceedings of the SIGGRAPH '87 Conference* 25 - 34.

Rose, C. 1999. Verbs and Adverbs: Multidimensional Motion Interpolation Using Radial Basis Functions. Ph.D. Dissertation, Department of Computer Science, Princeton University.

Schenkel, R. 1967. Submission: its features and function in the wolf and dog. *American. Zoologist*. 7:319-329.

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

teBoekhorst, I. and Hogeweg, P. 1994. Self-structuring in artificial 'CHIMPS' offers new hypothesis for male grouping in chimpanzees. *Behaviour* 130:229-52.

Tu, X. and Terzopoulos, D. 1994. Artificial Fishes: Physics, Locomotion, Perception, Behavior. In *Proceedings of ACM Computer Graphics, SIGGRAPH'94.* 43-50.

Velasquez, J. 1998. When Robots Weep: Emotional Memories and Decision-Making. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*. Madison, Wisconsin: AAAI Press.

Yoon, S., Blumberg, B. and Schneider, G. 2000. Motivation Driven Learning for Interactive Synthetic Characters. In *Proceedings of 4th Int'l. Conf. on Autonomous Agents (Agents 2000).*

Zimen, E. 1981. *The Wolf*. New York, NY: Delacorte Press