

# Synthetic Social Relationships for Computational Entities

by

William Michael Tomlinson, Jr.

Bachelor of Arts, Biology, Harvard College, 1994

Master of Fine Arts, Experimental Animation, California Institute of the Arts, 1996

Master of Science, Media Arts and Sciences, Massachusetts Institute of Technology, 1999

Submitted to the Program in Media Arts and Sciences,  
School of Architecture and Planning,  
in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy  
at the  
Massachusetts Institute of Technology  
June 2002

© 2002 Massachusetts Institute of Technology. All rights reserved.

---

Author

May 6, 2002

---

Certified by

**Bruce M. Blumberg**  
Associate Professor of Media Arts and Sciences  
Asahi Broadcasting Corporation Career Development Professor  
of Media Arts and Sciences  
Thesis Supervisor

---

Accepted by

**Andrew B. Lippman**  
Chairperson, Departmental Committee on Graduate Students



# Synthetic Social Relationships for Computational Entities

by  
William Michael Tomlinson, Jr.

Submitted to the Program in Media Arts and Sciences,  
School of Architecture and Planning,  
on May 6, 2002  
in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy  
at the  
Massachusetts Institute of Technology

## Abstract

Humans and many other animals form long term social relationships with each other. These relationships confer a variety of benefits upon us, both as individuals and as groups. Computational systems that can form social relationships like those formed by animals could reap many of the benefits of sociality, both within their own groups and in their interactions with people.

This dissertation explores two main questions:

- What kinds of internal and external representations are necessary for computational entities to form social relationships like those formed by animals?
- How can people participate in and direct the relationships of these entities?

To explore these questions, I designed and implemented a system by which computational entities may form simple social relationships. In particular, these synthetic social relationships are modeled after the social behavior of the gray wolf (*Canis lupus*). The system comprises a novel combination of simple models of emotion, perception and learning in an emotional memory-based mechanism for social relationship formation. The system also includes supporting technologies through which people may participate in and direct the relationships.

The system was presented as an interactive installation entitled AlphaWolf in the Emerging Technologies program at SIGGRAPH 2001. This installation featured a pack of six virtual wolves – three fully autonomous adults and three semi-autonomous pups whom people could direct by howling, growling, whining or barking into microphones.

In addition to observing the interactions of several hundred SIGGRAPH participants, I performed two main evaluations of the AlphaWolf system – a 32-subject human user study and a set of simulations of resource exploitation among the virtual wolves. Results from these evaluations support the hypothesis that the AlphaWolf system enables the formation of social relationships among groups of computational entities and people, and that these relationships are beneficial to both the inter-machine interactions and the human-machine interactions in a variety of ways.

This research represents one of many possible steps towards synthetic social relationships with the complexity of the relationships found in real wolves, let alone in humans. Much further work will be necessary to create entities who can fully engage us in our own social terms. The system presented here provides a basic scaffolding on which such entities may be built, including an implemented, real-time example; new ideas in directable characters and character-based interactive installations; a simple, ethologically plausible model of computational social relationships; and statistically significant support for these claims.

Thesis Supervisor: Bruce M. Blumberg

Title: Associate Professor of Media Arts and Sciences



## **Doctoral Dissertation Committee**

---

**Thesis Advisor**

**Bruce M. Blumberg, Ph.D.**  
**Associate Professor of Media Arts and Sciences**  
**Asahi Broadcasting Corporation Career Development Professor**  
**of Media Arts and Sciences**  
**M.I.T. Program in Media Arts and Sciences**

---

**Thesis Reader**

**Cynthia L. Breazeal, Sc.D.**  
**Assistant Professor of Media Arts and Sciences**  
**LG Career Development Professor of Media Arts and Sciences**  
**M.I.T. Program in Media Arts and Sciences**

---

**Thesis Reader**

**Rosalind W. Picard, Sc.D.**  
**Associate Professor of Media Arts and Sciences**  
**M.I.T. Program in Media Arts and Sciences**

---

**Thesis Reader**

**Richard W. Wrangham, Ph.D.**  
**Ruth B. Moore Professor of Anthropology**  
**Harvard University**

**May 2002**



## Acknowledgements

This thesis is largely about relationships among computational entities, but it has been relationships among people that have made it possible.

My first thanks go to my advisor Bruce Blumberg, who has kept me on track for the last five years, and who created the Synthetic Characters Group, my home during that time. It's thanks to you that I first became excited about our research, and thanks to you that I'm still excited about it five years later.

My three readers, Cynthia Breazeal, Roz Picard and Richard Wrangham, have been extremely thoughtful and thorough in supporting and directing this thesis. Your help in framing the research questions, evaluating the system and situating the work in a larger context has been invaluable. It has been a true pleasure to interact with each of you.

This thesis wouldn't have much to show for itself without the core team who created the AlphaWolf installation in the summer of 2001 – Marc Downie, Matt Berlin, Jesse Gray, Adolph Wong, Derek Lyons, Jennie Cochran and Bryan Yong. It was a great summer, and I'm proud to have been a part of the group who made AlphaWolf.

Mike Johnson, Marc Downie and Matt Berlin are the three other old-timers of the Synthetic Characters Group. Your friendship and advice has meant a lot to me over the last half-decade.

The first round of the Synthetic Characters Group – Mike Hlavac, Chris Kline, Ken Russell and Jed Wahl – showed me how to do the late-night installation-building that has become the Group's trademark. Dan Stiehl, my first real UROP and now one of the long-timers here at the Lab, has been a source of enthusiasm since we first made chickens in 1998. The rest of the Characters – Yuri Ivanov, Ben Resner, Rob Burke, Damian Isla, Scott Eaton, Andy Wilson, Song-Yee Yoon, Geoff Beatty, Delphine Nain, Rusmin Soetjipto, Dan Zaharopol, Matt Grimes and Lily Shirvanee – have provided companionship at all hours.

Aileen Kawabe has been instrumental in the AlphaWolf installation and throughout this thesis, both logistically and psychologically. Thanks for keeping me from falling apart at the seams!

Linda Peterson and Pat Solakoff helped keep me moving across the Big Graduate Student Board, and were nice to me even at the last minute.

Glorianna Davenport, my first real contact at the Media Lab, and Tom Porter advised me on my Master's thesis about expressive interactive cinematography. Their support of that project helped me figure out that relationships (which lie at the heart of cinematography) should be the topic for my doctoral research.

Thanks to all of the Media Lab crew at One Cambridge Center, in particular Ari Benbasat, Lisa Lieberman, Brygg Ullmer, Josh Lifton, Josh Strickon, Chi Yuen, Hiroshi Ishii and Joe Paradiso, and to everyone at the Media Lab mothership, past and present – Chris Wren, Deb Roy, Tony Jebara, Alan Wexelblat, Lenny Foner, Nelson Minar, Push Singh, Thad Starner, Robin Wilson, Marie Kerr, Alexandra Kahn, Ellen Hoffman, Deb Cohen, Matt Tragert, Henry Holtzman, and all the other people who kept the Lab running even when I didn't notice. Special thanks to my old roommate Brad Rhodes, too, who showed me how to do one of these Ph.D. things.

My beautiful girlfriend Marissa has been incredibly caring and supportive of me working all the time for the last couple years. You've been a model of rigor (and fun, too), and continue to be an inspiration for me!

Finally, I'd like to thank my family – Ma, Da, Lynn, Craig, Sam, Lucy, Uncle Bill, Grandfather Bill, Grammy Madaline, Grandfather Wesley, Grammy Joyce, Boomer, Petie and Pippin. There is some of each of you in the next 200 pages.

## Table of Contents

<b><u>Abstract</u></b> .....	<b>3</b>
<b><u>Doctoral Dissertation Committee</u></b> .....	<b>5</b>
<b><u>Acknowledgements</u></b> .....	<b>7</b>
<b><u>Table of Contents</u></b> .....	<b>9</b>
<b><u>List of Figures</u></b> .....	<b>15</b>
<b><u>1 Introduction</u></b> .....	<b>18</b>
1.1 <u>Motivation</u> .....	19
1.2 <u>Problem Statements</u> .....	20
1.3 <u>Contributions</u> .....	21
1.3.1 <u>Implemented, Real-time Example</u> .....	21
1.3.2 <u>Ideas in Directable Characters and Character-based Installations</u> .....	21
1.3.3 <u>Statistically Significant Verification</u> .....	21
1.3.4 <u>Simple, Ethologically plausible Model</u> .....	21
1.4 <u>Background</u> .....	21
1.5 <u>Key Terms</u> .....	22
1.6 <u>Approach</u> .....	23
1.6.1 <u>Computational Model</u> .....	24
1.6.2 <u>Directability</u> .....	25
1.7 <u>Intellectual Framework</u> .....	26
1.8 <u>Scope</u> .....	27
1.9 <u>Summary</u> .....	28
<b><u>2 The Alpha Wolf Installation</u></b> .....	<b>30</b>
2.1 <u>Sample Interaction</u> .....	31
2.2 <u>Emotibuttons</u> .....	33
2.3 <u>Graphics</u> .....	33
2.4 <u>Sound</u> .....	34
2.5 <u>Development</u> .....	34
2.6 <u>Cinematography</u> .....	35
2.7 <u>A Networked Pack</u> .....	35
2.8 <u>The Physical Set</u> .....	35
2.9 <u>The Audience</u> .....	36
2.10 <u>The Video</u> .....	36
2.11 <u>The Pitch</u> .....	36
2.12 <u>The Team</u> .....	37
<b><u>3 Background and Related Work</u></b> .....	<b>38</b>
3.1 <u>Natural Systems</u> .....	38
3.1.1 <u>The Gray Wolf</u> .....	38
3.1.1.1 <u>Life Cycle</u> .....	38
3.1.1.2 <u>Social Relationships</u> .....	40
3.1.1.3 <u>Research Efforts</u> .....	41
3.1.2 <u>Social Relationships</u> .....	43
3.1.3 <u>Emotion</u> .....	45
3.1.4 <u>Perception and Communication</u> .....	45
3.1.5 <u>Learning</u> .....	46

3.1.5.1	<a href="#">Emotional Learning</a>	47
3.1.5.2	<a href="#">Social Learning</a>	47
3.1.5.3	<a href="#">Training</a>	48
3.1.6	<a href="#">Development</a>	48
3.2	<a href="#">Computational Systems</a>	49
3.2.1	<a href="#">Synthetic Characters</a>	49
3.2.1.1	<a href="#">Action Selection</a>	49
3.2.1.2	<a href="#">Expression</a>	49
3.2.1.3	<a href="#">Learning</a>	49
3.2.1.4	<a href="#">Installations</a>	50
3.2.1.5	<a href="#">Other Systems</a>	51
3.2.2	<a href="#">Social Robotics</a>	52
3.2.3	<a href="#">Game-Theoretic Models</a>	52
3.2.4	<a href="#">Multi-Agent and Multi-Robot Systems</a>	53
3.2.5	<a href="#">Perception</a>	54
3.2.6	<a href="#">Emotion</a>	55
3.2.7	<a href="#">Learning</a>	55
3.2.8	<a href="#">Development</a>	56
3.2.9	<a href="#">Social Human-Computer Interaction</a>	56
3.2.10	<a href="#">Directable Characters and Agents</a>	57
3.2.11	<a href="#">Computer Games</a>	58
3.3	<a href="#">Summary</a>	59
<b>4</b>	<b><a href="#">Implementation</a></b>	<b>60</b>
4.1	<a href="#">Hardware</a>	60
4.2	<a href="#">The World</a>	61
4.2.1	<a href="#">The Three Implementations</a>	62
4.3	<a href="#">The Wolves</a>	62
4.4	<a href="#">Action Selection</a>	63
4.4.1	<a href="#">Visualizer</a>	66
4.4.2	<a href="#">The Three Implementations</a>	67
4.5	<a href="#">Navigation, Orientation and Collision</a>	67
4.5.1	<a href="#">Visualizer</a>	68
4.5.2	<a href="#">The Three Implementations</a>	68
4.6	<a href="#">Emotion and Motivation</a>	69
4.6.1	<a href="#">Visualizer</a>	70
4.6.2	<a href="#">The Three Implementations</a>	70
4.7	<a href="#">Sensation and Perception</a>	70
4.7.1	<a href="#">To Whom To Attend</a>	73
4.7.2	<a href="#">Look At</a>	73
4.7.3	<a href="#">Visualizer</a>	75
4.7.4	<a href="#">The Three Implementations</a>	76
4.8	<a href="#">Context Specific Emotional Memories (CSEMs)</a>	76
4.8.1	<a href="#">Visualizers</a>	78
4.8.1.1	<a href="#">Instantaneous</a>	78
4.8.1.2	<a href="#">Continuous</a>	79
4.8.2	<a href="#">An Example</a>	80
4.8.3	<a href="#">The Three Implementations</a>	82
4.9	<a href="#">Development</a>	83
4.9.1	<a href="#">Physical Development</a>	83
4.9.2	<a href="#">Behavioral Development</a>	83

4.9.3	<a href="#">Visualizer</a>	83
4.9.4	<a href="#">The Three Implementations</a>	83
4.10	<a href="#">Expression</a>	84
4.10.1	<a href="#">Animation</a>	84
4.10.2	<a href="#">Sound</a>	85
4.10.3	<a href="#">Visualizer</a>	85
4.10.4	<a href="#">The Three Implementations</a>	85
4.11	<a href="#">Cinematography</a>	86
4.11.1	<a href="#">Visualizer</a>	87
4.11.2	<a href="#">The Three Implementations</a>	87
4.12	<a href="#">Interface</a>	88
4.12.1	<a href="#">Microphone</a>	88
4.12.2	<a href="#">Mouse</a>	88
4.12.3	<a href="#">Visualizer</a>	88
4.12.4	<a href="#">The Three Implementations</a>	89
4.13	<a href="#">Summary</a>	89
<b>5</b>	<b><a href="#">Evaluation</a></b>	<b>90</b>
5.1	<a href="#">Summary of Results</a>	90
5.1.1	<a href="#">Transmission of Relationships</a>	90
5.1.2	<a href="#">Effect of Social Relationship Algorithms</a>	90
5.1.3	<a href="#">Effect of Interaction Paradigms</a>	91
5.1.4	<a href="#">Simulations of Resource Exploitation</a>	92
5.1.5	<a href="#">The AlphaWolf Installation at SIGGRAPH</a>	92
5.2	<a href="#">User Tests</a>	92
5.2.1	<a href="#">Hypotheses</a>	94
5.2.2	<a href="#">Experimental Method</a>	94
5.2.2.1	<a href="#">Alternate Algorithms</a>	94
Emotional		94
Fixed		94
Random		95
5.2.2.2	<a href="#">Interaction Paradigms</a>	95
Microphone		95
Button		95
Non-interactive		95
5.2.2.3	<a href="#">Procedure</a>	96
Preparation		96
Introduction		96
Video of Real Wolves		97
The Three Runs		97
The Three Questionnaires		98
Final Questionnaire		99
Debrief		99
Process Data		99
5.2.2.4	<a href="#">Subjects</a>	99
Gender Effects		100
Age Effects		100
5.2.2.5	<a href="#">The Virtual Wolves</a>	100
5.2.3	<a href="#">Results</a>	101
5.2.3.1	<a href="#">Encoding, Decoding and Transmission</a>	101
Shannon's Theory of Communication		101

<u>Information Source</u>	102
<u>Transmitter</u>	102
<u>Channel</u>	102
<u>Receiver</u>	102
<u>Destination</u>	102
<u>Encoding</u>	103
<u>Decoding</u>	104
<u>The Relationships between Black and White</u>	105
<u>Transmission</u>	106
<u>Discussion</u>	107
<u>Black and White</u>	108
5.2.3.2 <u>Effects of Social Relationship Algorithms</u>	108
<u>Experimental Method</u>	109
<u>Results</u>	109
<u>Discussion</u>	110
<u>A Preference for Emotion?</u>	111
5.2.3.3 <u>Interaction Paradigms</u>	112
<u>Experimental Method</u>	112
<u>Results</u>	112
<u>Similarity to Real Wolves</u>	113
<u>Clarity of Relationships</u>	114
<u>User Enjoyment</u>	115
<u>Control</u>	115
<u>Comfort</u>	117
<u>Liking of Wolves</u>	117
<u>Identification with Wolves</u>	119
<u>Immersion</u>	121
<u>Discussion</u>	122
5.2.3.4 <u>Additional Thoughts</u>	123
<u>Pronouns and the Relationship between Subject and Pup</u>	123
<u>Telling Stories</u>	124
<u>Enthusiasm and Learning</u>	124
<u>Expressiveness</u>	124
<u>Problems</u>	124
<u>Close Contact vs. Interpenetration</u>	124
<u>Sound Effects</u>	125
<u>Camera</u>	125
5.3 <u>Simulations</u>	126
5.3.1 <u>AlphaWolves Form Hierarchies</u>	127
5.3.1.1 <u>Other Algorithms</u>	129
5.3.2 <u>Hierarchies Create an Imbalance in Resource Distribution</u>	131
5.3.3 <u>AlphaWolf Hierarchies Are Directable</u>	134
5.3.4 <u>Discussion</u>	136
5.4 <u>The AlphaWolf Installation</u>	138
5.4.1.1 <u>Press</u>	139
<b>6 <u>Discussion</u></b>	<b>141</b>
6.1 <u>Directable Social Characters</u>	141
6.1.1 <u>Necessary Components</u>	141
6.1.2 <u>Clear Emotions</u>	142
6.1.3 <u>Controllable Action</u>	143



6.1.4	<a href="#">Directability</a>	144
6.2	<a href="#">Character-Based Interactive Installations</a>	145
6.2.1	<a href="#">Expectation Management</a>	145
6.2.2	<a href="#">Game Play</a>	146
6.2.3	<a href="#">Interactivity</a>	147
6.2.4	<a href="#">Social Is Intuitive</a>	148
6.2.5	<a href="#">Social Is Effective</a>	148
6.2.6	<a href="#">The Implications of Dominance</a>	149
6.2.7	<a href="#">Implicit Biases</a>	149
6.2.8	<a href="#">Simplicity</a>	150
6.2.9	<a href="#">Real Time</a>	151
6.2.10	<a href="#">Stories</a>	151
6.2.11	<a href="#">Cinematography</a>	151
6.2.12	<a href="#">Sound</a>	152
6.2.13	<a href="#">Physical and Contextual Setting</a>	153
6.2.14	<a href="#">Process</a>	153
6.3	<a href="#">Three Kinds of Relationships</a>	154
6.3.1	<a href="#">Relationships between Machines and Machines</a>	154
6.3.2	<a href="#">Relationships between Humans and Machines</a>	155
6.3.2.1	<a href="#">Relationship between Participant and Gray Pup</a>	155
6.3.2.2	<a href="#">Relationship between Participant and Other Wolves</a>	157
6.3.3	<a href="#">Relationships between Humans and Humans</a>	158
<b>7</b>	<b><a href="#">Future Work and Applications</a></b>	<b>159</b>
7.1	<a href="#">Limitations</a>	159
7.1.1	<a href="#">Emotional Simplicity</a>	159
7.1.2	<a href="#">Perceptual Simplicity</a>	159
7.1.3	<a href="#">Behavioral Simplicity</a>	159
7.1.4	<a href="#">Relationship Bandwidth</a>	159
7.1.5	<a href="#">No Memories of Specific Events</a>	160
7.1.6	<a href="#">Real-time</a>	161
7.1.7	<a href="#">Supporting Technology Problems</a>	161
7.1.8	<a href="#">Novelty Can Be a Problem</a>	161
7.2	<a href="#">Future Work</a>	162
7.2.1	<a href="#">Scaling</a>	162
7.2.2	<a href="#">Single-Sex Hierarchies</a>	163
7.2.3	<a href="#">Deposed Alpha Plummets</a>	163
7.2.4	<a href="#">Coopting Behaviors to Serve as Social Signals</a>	164
7.2.5	<a href="#">Facilitating Social Learning</a>	164
7.2.6	<a href="#">Signaling, Deceiving and Detecting Deception</a>	164
7.2.7	<a href="#">Alliance Formation</a>	165
7.2.8	<a href="#">Conjunctions</a>	165
7.2.9	<a href="#">More Complex Emotional Models</a>	166
7.2.10	<a href="#">Formalization of Development</a>	166
7.2.11	<a href="#">Synthetic Sensing for Assessing Dominance</a>	166
7.2.12	<a href="#">Close Contact</a>	167
7.3	<a href="#">Applications</a>	168
7.3.1	<a href="#">Entertainment</a>	168
7.3.1.1	<a href="#">Computer Games</a>	168
7.3.1.2	<a href="#">Television and Movies</a>	168
7.3.1.3	<a href="#">Toys</a>	169

7.3.1.4	<a href="#">Location-Based Entertainment</a>	169
7.3.1.5	<a href="#">New Forms</a>	169
7.3.2	<a href="#">Education</a>	170
7.3.3	<a href="#">Robotics</a>	171
7.3.4	<a href="#">E-business</a>	172
7.3.5	<a href="#">Synthetic Biology and Artificial Life</a>	173
7.3.6	<a href="#">Fine Art</a>	173
7.3.7	<a href="#">Computer-Mediated Communication Technologies</a>	174
7.3.8	<a href="#">Human Computer Interaction</a>	174
7.3.9	<a href="#">Artificial Intelligence</a>	175
<b>8</b>	<b><a href="#">Conclusions</a></b>	<b>176</b>
	<b><a href="#">Bibliography</a></b>	<b>179</b>
	<b><a href="#">Appendix: Experimental Protocol</a></b>	<b>190</b>
	<a href="#">Experimenter's Script</a>	190
	<a href="#">Consent Form</a>	193
	<a href="#">Questionnaire</a>	194
	<a href="#">Final Questionnaire – Microphone/Button/Video</a>	196
	<a href="#">Final Questionnaire – Button 1/Button 2/Video</a>	198
	<a href="#">Debriefing Statement</a>	200

## List of Figures

<a href="#"><u>Figure 1-1: Two social computational entities exchange a glance.</u></a>	18
<a href="#"><u>Figure 1-2: The three kinds of relationships enabled by a computational mechanism for social relationship formation. 1) Human-machine. 2) Machine-machine. 3) Human-human, mediated by machines.</u></a>	20
<a href="#"><u>Figure 1-3: Wolf social behavior in the wild and in the AlphaWolf installation.</u></a>	24
<a href="#"><u>Figure 1-4: A gray wolf pup looks at its sleeping sibling.</u></a>	28
<a href="#"><u>Figure 2-1: The AlphaWolf logo.</u></a>	30
<a href="#"><u>Figure 2-2: A participant (Michael Wahrman) directs the actions of his wolf.</u></a>	30
<a href="#"><u>Figure 2-3: The full AlphaWolf installation, with three kiosks and a plasma display.</u></a>	31
<a href="#"><u>Figure 2-4: The pups start off asleep.</u></a>	31
<a href="#"><u>Figure 2-5: Black dominates white.</u></a>	32
<a href="#"><u>Figure 2-6: The three pups, grown up now, have a final howl together.</u></a>	34
<a href="#"><u>Figure 2-7: A pup tries to wake up an adult in the AlphaWolf pitch, January 2001.</u></a>	36
<a href="#"><u>Figure 2-8: Some Synthetic Characters make some synthetic characters.</u></a>	37
<a href="#"><u>Figure 3-1: Two gray wolves, one gray in color, one black, having a dominance interaction.</u></a>	38
<a href="#"><u>Figure 3-2: Developmental stages of the gray wolf (not to scale).</u></a>	39
<a href="#"><u>Figure 3-3: Timeline of an average gray wolf's life (not to scale).</u></a>	39
<a href="#"><u>Figure 3-4: Passive submission (from [Schenkel 1967]).</u></a>	40
<a href="#"><u>Figure 3-5: Active submission (from [Schenkel 1967]).</u></a>	41
<a href="#"><u>Figure 3-6: Active submission in adults resembles food soliciting behavior in pups.</u></a>	42
<a href="#"><u>Figure 3-7: The author with a captive gray wolf at Wolf Park in Battle Ground, Indiana.</u></a>	43
<a href="#"><u>Figure 3-8: A dimensional emotional model.</u></a>	45
<a href="#"><u>Figure 3-9: A wolf communicates.</u></a>	46
<a href="#"><u>Figure 3-10: The raccoon in Swamped! at SIGGRAPH '98.</u></a>	50
<a href="#"><u>Figure 3-11: Elliott the Salesman from (void*).</u></a>	50
<a href="#"><u>Figure 3-12: Duncan herds ornery sheep in sheep dog.</u></a>	51
<a href="#"><u>Figure 3-13: Kismet.</u></a>	52
<a href="#"><u>Figure 3-14: Craig Reynolds' Boids (from <a href="http://www.red3d.com/cwr/boids/">http://www.red3d.com/cwr/boids/</a>).</u></a>	53
<a href="#"><u>Figure 3-15: A Norn from the game Creatures.</u></a>	56
<a href="#"><u>Figure 3-16: <i>The Sims</i>.</u></a>	58
<a href="#"><u>Figure 4-1: The hardware behind the full AlphaWolf installation.</u></a>	61
<a href="#"><u>Figure 4-2: The gray wolf walks around the virtual forest.</u></a>	62
<a href="#"><u>Figure 4-3: The pup's skeleton and skin.</u></a>	63
<a href="#"><u>Figure 4-4: An ActionTuple.</u></a>	63
<a href="#"><u>Figure 4-5: The Action System visualizer for a wolf in the act of whining.</u></a>	66
<a href="#"><u>Figure 4-6: These visualizers show that Wolf5 has a very low dominance and a mid-range hunger.</u></a>	70
<a href="#"><u>Figure 4-7: A sample interaction as perceived by the Wolf B.</u></a>	72
<a href="#"><u>Figure 4-8: When the wolf's object of attention is within a certain angle of its forward vector, it attends. If it is beyond this angle, it turns.</u></a>	74
<a href="#"><u>Figure 4-9: The gray wolf pup showing a strong look at skill.</u></a>	74
<a href="#"><u>Figure 4-10: The Percept Tree.</u></a>	76
<a href="#"><u>Figure 4-11: The visualizer of all the pups' relationships.</u></a>	78
<a href="#"><u>Figure 4-12: A continuous representation of the pups' dominances.</u></a>	80
<a href="#"><u>Figure 4-13: The first 1000 virtual seconds in a four-member pack.</u></a>	80
<a href="#"><u>Figure 4-14: t = 50. The pups are just getting to know each other.</u></a>	81
<a href="#"><u>Figure 4-15: t = 100. All the pups have met each other. A few of the relationships have settled in.</u></a>	81

Figure 4-16: $t = 400$ . All of the relationships have stabilized except the one between white and black.	82
Figure 4-17: $t = 1000$ . The dominance relationships are now well established, with gray at the top, then white, then black, and red at the bottom.	82
Figure 4-18: An emotional range from submissive to dominant. Our animator crafted the extreme poses. The center three poses are automatically generated.	84
Figure 4-19: The camera angle that guides people back to the center of the world.	87
Figure 5-1: Diagram of human user study's experimental method.	93
Figure 5-2: The experimental setup.	96
Figure 5-3: Schematic diagram of a general communication system (from [Shannon 1948]).	101
Figure 5-4: Subject A is assigned the task of forming two social relationships. He or she encodes these relationships into a computational representation through an interactive run of the AlphaWolf system. These relationships are loaded in at the beginning of Subject B's Non-interactive run. If Subject B can perceive the relationships, they have been successfully decoded. If both encoding and decoding work, transmission occurs.	102
Figure 5-5: The AlphaWolf mechanism offers significantly improved encoding over random chance.	104
Figure 5-6: Decoding is successful regardless of encoding mechanism.	105
Figure 5-7: Transmission is significantly better with the AlphaWolf mechanism than random chance.	107
Figure 5-8: The AlphaWolf mechanism offers significantly better encoding, and therefore transmission, than chance (50%) or than a system without the AlphaWolf mechanism. Subjects were successful at decoding relationships regardless of how they were encoded.	108
Figure 5-9: Subjects preferred the second Button Run, regardless of algorithm.	110
Figure 5-10: AlphaWolf was favored with more responses than the other algorithms were when a first Button run was chosen over a second Button run.	110
Figure 5-11: Subjects found the pups in Non-interactive runs to be more like real wolves than those in interactive runs.	113
Figure 5-12: Relationships are clearer in Non-interactive runs.	114
Figure 5-13: Subjects felt more control in the interactive runs.	116
Figure 5-14: Subjects felt slightly more control with buttons than with the microphone.	116
Figure 5-15: Cinematography makes more of a difference than interactivity in how much people like pups.	117
Figure 5-16: People tend to like the wolves more in later Non-interactive runs than earlier ones.	118
Figure 5-17: Cinematography matters more than interactivity in identification, too.	120
Figure 5-18: People identify more with the gray pup in non-interactive runs, even when it's their first run.	120
Figure 5-19: Subjects identify more with all three pups when interacting with a microphone.	121
Figure 5-20: Interactive experiences with the wolves are more immersing than non-interactive experiences.	122
Figure 5-21: The first 500 virtual seconds in a newly created pack.	127
Figure 5-22: The same pack as in the previous figure, now 8000 seconds old. The dominance relationships have largely stabilized.	128
Figure 5-23: The same pack, with crossing-over among middle-ranking wolves.	129
Figure 5-24: The first 1000 seconds of a pack featuring the Emotional algorithm.	130
Figure 5-25: The first 1000 seconds of a pack with the Fixed algorithm.	130
Figure 5-26: The first 1000 seconds of a pack with Random social relationships.	131
Figure 5-27: A dominant wolf can feed whenever it wants (here, when hunger reaches 0.7).	132
Figure 5-28: A submissive wolf often gets interrupted when feeding, resulting in a much more jagged hunger profile.	133

<a href="#">Figure 5-29: The Fixed algorithms creates a pack with strong polarization of hunger levels. AlphaWolf is next, followed by Emotional. The Random algorithm shows very little polarization of hunger among the pack members.</a>	134
<a href="#">Figure 5-30: A pack of entities with the AlphaWolf mechanism are directable. Every 4000 seconds, a wolf is chosen for “promotion”, and is forced to be dominant in all its interactions for 600 seconds. The effect is most evident among the promotions of lower-ranking wolves (Red at t = 12000, White at t = 16000, Gray at t = 20000).</a>	135
<a href="#">Figure 5-31: When a wolf is promoted, its hunger profile becomes that of a dominant wolf.</a>	136
<a href="#">Figure 5-32: Each pup is asleep at the beginning of the interaction.</a>	138
<a href="#">Figure 6-1: The gray pup walking sadly.</a>	142
<a href="#">Figure 6-2: The flow of control among a user and two pups.</a>	144
<a href="#">Figure 6-3: People like and identify with their gray pup.</a>	152
<a href="#">Figure 7-1: Limitations on social bandwidth are imposed by simplification and interfaces. The interaction bandwidth between the two people, for example, is greatly compressed by passing it through the virtual wolves (not drawn to scale).</a>	160
<a href="#">Figure 7-2: Close contact – one of the reasons that social behavior is a hard problem.</a>	167
<a href="#">Figure 7-3: The author, about to awaken a sleeping virtual wolf pup.</a>	174
<a href="#">Figure 8-1: The black wolf howls to find its pack mates.</a>	177

# 1 Introduction

Humans and many other animals form long term social relationships with each other. These relationships confer certain benefits upon us, both individually and as groups. For example, long-term dyadic social relationships enable the existence of the stable alliances that we see in primate societies [Harcourt 1992; Barrett 2000]. Pair-bonding, among wolves for example [Fox 1971], relies upon a persistent relationship between two individuals, and in the case of wolves provides for at least one other individual to hunt while the mother is in the den with newborn pups. The ability to preserve a different behavioral context [Cohen 1999] with each of our social partners allows people to interact with, cooperate with, and learn from each member of our social group in a consistently different way. The value of social relationships is all around us, in our diverse societies and cultures.

Currently, computer systems are not able to participate in social relationships of the kind that people and animals form. They are neither able to harness the full benefits of sociality among their own groups, nor are they able to engage people with the richness that we expect from natural social entities (although significant efforts have certainly been made in both of these areas, which will be discussed in the Related Work chapter of this document).

The goal of this thesis has been to create a simple system through which computational entities may form social relationships that resemble those formed by people and animals. In addition, the system endeavors to make it possible for people to participate in and direct those relationships. A simple social relationship mechanism for computational entities could help form the basis for a wide variety of applications in computing and its interactions with humanity. For example, a group of autonomous robots might use social relationships to coordinate their actions. Or the members of a retirement community might be able to make friends with a new synthetic pet, each in his or her own way. Or a child might be able to learn about wolf social behavior by getting a chance to be a member of a simulated pack. Or a movie director might be able to say “Action” and an autonomous computational world would spin into action, the characters interacting with each other in socially plausible ways.



Figure 1-1: Two social computational entities exchange a glance.

A social relationship has several clear attributes. It is learned during an individual's lifetime, rather than genetically "hard-coded." It changes continually during this time based on the ongoing interaction history between the individual and its social partner. It is closely linked with emotion, both in its formation [Damasio 1994] and in its expression [Darwin 1965 (originally published 1872)]. The computational model presented in this document accounts for each of these attributes, incorporating simple models of emotion, perception and learning. This thesis presents a novel mechanism for the formation and maintenance of multiple synthetic social relationships, and is the first time that an emotion-based social relationship mechanism has been applied in an interactive domain, thereby allowing people to participate in the relationships.

In order to have a concrete target for this undertaking, I have chosen a specific natural example, the gray wolf (*Canis lupus*), as the focus for a working implementation. Canids (wolves and dogs) form long-term dyadic social relationships that are clearly visible through their behavior [Mech 1998]. In addition, these relationships interface readily with human social relationships (as any dog-owner knows). In order to demonstrate that I have created a system through which computational entities can form social relationships with each other, I have created an interactive 3D-animated simulation of a wild wolf pack (see Figure 1-1). An installation based on the virtual wolf pack premiered under the title AlphaWolf in the Emerging Technologies program at SIGGRAPH 2001 [Tomlinson 2001].

The synthetic social relationships presented in this thesis pale in comparison to those of real wolves, let alone to those that form the basis of human societies. Nevertheless, I hope that the ideas presented here help pave the way for more complex and interesting forms of social relationships for computational entities.

## **1.1 Motivation**

There are three main groups of reasons why the building of computational systems with the ability to form social relationships is an important undertaking.

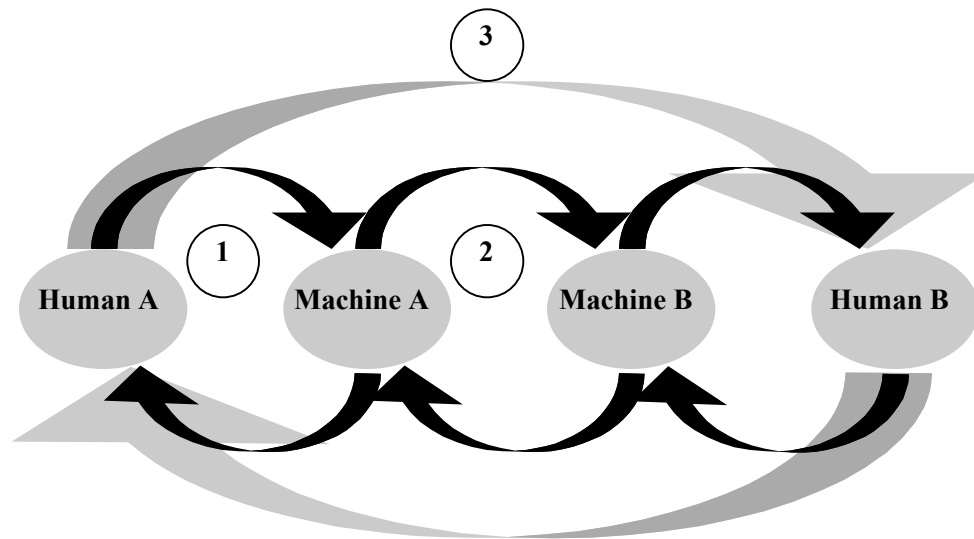
First, as Stanford researchers Byron Reeves and Clifford Nass have noted, humans try to apply our interpersonal social skills to the technology around us [Reeves 1996]. As most computer users have probably experienced at one point or another, it can be frustrating when the machines do not react in kind. While computational systems with social skills are not a panacea (I'm not suggesting that air traffic control computers should ever "have a bad day"), there are nevertheless a variety of possible applications for socially enabled computational systems in the arena of human-computer interaction. As an example, imagine a stuffed animal that could form relationships with the children who played with it, interacting with each one differently. Or a synthetic social partner for a person who would like a pet, but is physically unable to care for one.

Second, just as social relationships serve a purpose in non-human natural systems without reference to humanity (e.g., for resource allocation, breeding success, and cooperative hunting [Alcock 1989]), so too could social relationships increase the functionality of groups of computational entities. Multi-robot systems, for example, could use social relationships to coordinate the way in which they exploit the resources in their environments, allowing more "dominant" robots first access [Mataric 1995].

Finally, consider the many social relationships among humans that are now maintained through technology. Friends and family separated by distance use telephones, email and instant messaging to keep in touch; human social relationships are one of the main reasons that ordinary people embrace technology. Whenever technology intervenes in interpersonal communication, it

“mediates” it. Sometimes the medium is meant to be largely transparent (e.g., telephone static is a bad thing), and sometimes it is intended to mediate the interaction actively (e.g., avatar-based chat rooms, intelligent conference rooms [Jebara 2000]). If the mediating technology had some cognizance of human social relationships, it might be able to do its job better. As an example of how this kind of functionality could be applied, imagine a massively multiplayer networked computer game, where each player’s avatar has an idea of what kind of body language to use based on the user’s interaction history with the other avatars with whom it is interacting.

The three areas that I just mentioned point to three distinct kinds of relationships that might be created by a computational social relationship mechanism – human-machine, machine-machine, and human-human. Each of these kinds of relationship is valuable in its own right; making machines that can form social relationships will enhance all three in various ways. Examples from these three domains of social relationship will be included throughout this document.



**Figure 1-2: The three kinds of relationships enabled by a computational mechanism for social relationship formation. 1) Human-machine. 2) Machine-machine. 3) Human-human, mediated by machines.**

## ***1.2 Problem Statements***

In order to create a system by which computational entities can form relationships with each other and with people, two main problems must be addressed:

- What kinds of internal and external representations are necessary for computational entities to form social relationships like those formed by animals?
- How can people participate in and direct the relationships of these entities?

These two questions are the central issues to be implemented, evaluated and discussed in this document. In particular, they will be applied to the specific area of wolf social behavior, and then extrapolated to a variety of other potential social systems.



## ***1.3 Contributions***

There are several central contributions made by this dissertation toward the above questions, described below.

### **1.3.1 Implemented, Real-time Example**

The AlphaWolf installation, described fully in Chapter 2, is a functional implementation of the ideas in this document. It runs in real-time, and has engaged approximately one thousand participants since August 2001.

### **1.3.2 Ideas in Directable Characters and Character-based Installations**

The process of creating the AlphaWolf system and watching participants interact with it has brought to light a number of new ideas in the design of directable characters and interactive installations based on those characters. These ideas are described throughout this thesis, and in particular in Chapter 6.

### **1.3.3 Statistically Significant Verification**

Several evaluations were conducted on the AlphaWolf system, verifying that the system does, in fact, form relationships among virtual characters that may be perceived by people, and that various factors in the installation design had a profound impact on the subjective experience of that installation. These evaluations are detailed in Chapter 5.

### **1.3.4 Simple, Ethologically plausible Model**

The mechanism that lies at the heart of the AlphaWolf system is novel, and effectively simulates a subset of wolf social behavior. This mechanism and the system in which it is embedded are described in Chapter 4.

## ***1.4 Background***

The research project that I have undertaken to explore the above questions is part of a group effort under way in the Synthetic Characters Group at the MIT Media Lab, under the direction of Professor Bruce Blumberg. Our collective goal is to make computational entities that can do some of the things that animals do. Nature does a great job of doing apparently simple things like walking, seeing, learning, or growing (all of which many engineers will tell you are not that simple when you try to make machines do them). Great success has been had, historically, by taking the right lessons from nature. (Consider the Wright brothers, who made the realization that it was the shape of the wing, rather than the flapping, that makes controlled flight possible for birds). Ultimately we hope to help make computational creatures with the everyday common sense, ability to learn, and expressiveness of animals such as dogs. To this end, we have created a number of previous interactive installations to showcase our work (e.g., [Blumberg 1998; Blumberg 1999]). These and other previous projects are discussed in greater depth in the Related Work chapter of this document.

The main technical focus of our group is to create a single, unified tool kit for building virtual creatures. Each component needs to integrate cleanly with the others to be of universal use. The social relationship mechanism is therefore not just a single research project, but also a novel contribution to a larger, continuing effort. If this research in social relationship formation is successful in the long term, all of the Synthetic Characters' characters from here on in will have the ability to form social relationships.

In our work, the Synthetic Characters Group take a “whole-system approach.” We believe that it is not possible to understand or build part of a complex system without having that part situated in the larger context of the whole system in which it will be operating. This theme is prominent in certain modern views of artificial intelligence, for example, the notion of “embodied intelligence” [Steels 1995].

In particular, our version of a whole-system approach is a “character-based approach.” This means that we build “characters” – fully featured virtual entities that can interact with their environments in a variety of ways. Rather than working on just motor control or just navigation or just learning or a number of other topics, we work on all of these at once, in a hope that each will benefit from the synergy. The social relationship mechanism described in this thesis was made possible by its tight coupling with the myriad existing systems in our character-building tool kit (e.g., [Burke 2001; Downie 2001a; Isla 2001b; Blumberg 2001 (to appear)]), and will hopefully benefit them in return.

Despite the close ties between the work presented here and the other research efforts in our group, the social relationship mechanism is an original and significant research project in its own right. The distinctions between the project itself and the surrounding scaffolding is made clear in the Implementation chapter of this thesis. The importance of the social relationship research is demonstrated by the fact that it forms the core of the AlphaWolf installation, one of our group’s most recent large research projects. This project will be described in greater detail in Chapter 2.

## ***1.5 Key Terms***

In order to start from some common ground, I now offer a few definitions of key terms used throughout this thesis. Many of these topics will be addressed in much greater depth in the Related Work chapter and elsewhere. This section is just a quick introduction so that the terminology is clear.

### **Entity**

For the purposes of this research, an entity is defined as anything which can act in an autonomous fashion on its environment and react to that environment. This definition is fairly comprehensive, and should be seen as approximately equivalent to “autonomous actor” but without the computational overtones that “autonomous” has recently acquired. Specifically, an entity is expected to exhibit constrained unpredictability in its behavior as perceived by another entity. Modern computers are encompassed by this description; they exhibit autonomous behavior, and add constrained unpredictability to this behavior by crashing spontaneously, making suggestions about spelling mistakes, etc.

### **Social**

While the Oxford English Dictionary lists at least 12 major definitions of the term “social” [OED 2002], I will be relying primarily on definition 6b: “*Zool.* Living together in more or less organized communities; belonging to a community of this kind.” This definition is fairly clear, and lacks any of the human-imposed value judgments associated with being social. The social behavior that I have tried to create in the virtual wolves is social in the biological sense – “pertaining to group-living.” Nevertheless, some of the lessons learned apply to the various humanistic senses of the term “social.”

### **Social Relationship**

For the purposes of this research, a social relationship is a learned and remembered construct by which an entity keeps track of its interaction history with another entity, and allows that history to

affect its current and future interactions with that entity. Context preservation [Cohen 1999] is the essential function of a social relationship – behaving differently toward different social partners, rather than interacting with all entities in the same way.

The feature of this definition that makes it different from any other kind of learning (all of which could be encompassed by the above definition, depending on the exact definitions of terms like “entity” and “interaction history”), is that it pertains to other social entities. There is an expectation in a social relationship that there will be a “back-and-forth”, with both individuals providing some input to the interaction.

One of the most interesting things about social relationships is how they develop. Social groups are not static; rather, they continually have members added and removed, and the individual members continually change. In order to account for this constant flux, social relationships need to be adaptable and dynamic. As will be shown later in this thesis, the synthetic social relationships described here are able to change dynamically (e.g., two individuals reversing roles).

### **Dominance Relationship**

Of the varied kinds of social relationships that occur in the natural world, dominance relationships, particularly among wolves, tend to be strong and well-expressed. De Vries offers a functional definition of a dominance relationship, based on the “winning” of interactions. “A simple and straightforward way [to determine a dominance relationship], which is more or less standard in social dominance studies, is to call A dominant to B if A wins more encounters from B than B wins from A.” [de Vries 1998] In the case of wolves, winning would mean causing the other wolf to exhibit certain characteristic submission behaviors.

### **Computational Entity**

I will be using this term to mean any computational system that can be interacted with as an individual by a human or another computational entity. Examples include computer game characters, robots, avatars, and autonomous agents. It is my belief that there will be many more computational entities, in an assortment of guises, in the years to come. By choosing this general term, rather than a more specific term such as “agent”, I hope to encompass these future forms.

### **Synthetic Social Behavior**

In the spirit of Minsky [Minsky 1968], I define synthetic social behavior as the science of making machines do things that would be called social if done by people or animals.


### **Synthetic Social Relationship**

A synthetic social relationship is a computational construct by which one entity stores its interaction history with another entity and by which its behavior is affected. One particular instance of this class of possible representations is described in detail in this thesis. Nevertheless, many of the themes discussed here should be appropriate to a broader range of types of synthetic social relationships, regardless of the specific implementation.

## ***1.6 Approach***

In order to implement the synthetic social relationships described in this thesis, it was necessary to focus on one specific example that we could use to prove that it works. We chose wolves as the model for our simulation (see Figure 1-3) 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 our previous installations that have featured virtual dogs [Blumberg 1996; Burke 2001]. Nevertheless, wolves are sufficiently distinct from dogs that they provide a chance for us to extend and generalize our models. Finally, the social behaviors of wolves are similar enough to those of humans that some of the lessons we learn from the virtual wolves (e.g., that dominance hierarchies may be the natural result of simple emotional memories) might be relevant to human social behavior and simulation.



**Figure 1-3: Wolf social behavior in the wild and in the AlphaWolf installation.**

The research in this document is derived from existing literature on wolf behavior and social behavior (e.g., [Fox 1971; Klinghammer 1985; Mech 1998]), and also on other kinds of natural social relationships (e.g., [Wrangham 1994; Heyes 1996]). It draws on research about emotion (e.g., [Darwin 1965 (originally published 1872); Mehrabian 1974; Picard 1997]), neurology (e.g., [Damasio 1994]), and development (e.g., [Lorenz 1973] [Weng 2001]). In addition, this work has been done with an awareness of previous research projects that have dealt with computational entities with social abilities of one kind or another (e.g. [Reynolds 1987; Hemelrijk 1996; Breazeal 2000]). The mechanism described below is simple, effective and biologically plausible, and has not previously been implemented.

To demonstrate the necessary and sufficient elements required to build synthetic social behavior, let's look once again at the definition I offered above. I suggested that a social relationship is “a learned and remembered construct by which an entity keeps track of its interaction history with another entity, and allows that history to affect its current and future interactions with that entity.” What are the critical parts of that definition, with regard to a computational implementation?

- “learned and remembered construct” – there must be persistent storage of information about social partners, and that information must be acquired over the course of one or more interactions
- “entity” – there must be a notion of an entity
- “interaction history” – interactions must have an effect on an individual, and in turn on the remembered construct
- “another entity” – an entity must be aware of its social partners as entities that persist between interactions
- “affect its interactions” – the remembered construct must have an impact on behavior

### **1.6.1 Computational Model**

Based on the above definition, the social relationship mechanism described in this thesis involves simple models of emotion, perception, and learning. Each wolf maintains a continuously changing emotional state that is affected by its interactions (e.g., if another wolf dominates me, I

feel more submissive). A wolf can perceive the presence of other individuals in a repeatable fashion (e.g., it can recognize the white wolf as the same individual each time they meet). At the end of its first interaction with another wolf, that wolf forms an association (an “emotional memory”) between its current emotional state and the presence of that other wolf (who is now a “social partner” for it). When the wolf again encounters that social partner, the emotional memory influences its current emotional state to a degree determined by its confidence in that relationship, so that it can “pick up where it left off” with regard to its emotional relationship. At the end of each interaction, the wolf revises its emotional relationship with that social partner and its confidence in that relationship. This mechanism allows individual wolves to interact differently with specific social partners, based on the history of interactions between them.

This representation will be called a Context Specific Emotional Memory (CSEM). The “Emotional Memory” part of the name is clear enough – individuals form memories of the emotions that they tend to feel around certain social partners. “Context Specific” demonstrates that these Emotional Memories are attached to certain contextual information drawn from that wolf’s perceptual world, in this case, the perception of a certain individual in the wolf’s proximity. “Context Specific” was chosen instead of “Partner Specific” because, in a potential extension of the currently implemented system, different relationships could be formed with the same social partner under different surrounding conditions. For example, it could be possible for a wolf pup to have a different relationship with an older sibling when a parent is present versus when the parents aren’t around. This possible extension will be discussed further in the Future Work section of this document.

In addition to the term CSEM, the central mechanism described above will also be referred to as “the AlphaWolf mechanism.”

To perform the technical implementation of the AlphaWolf mechanism, I started with existing parts of the Synthetic Characters’ system (e.g., perception, navigation, motor control) and created a new system to draw these parts together. The specifics of the algorithms behind this mechanism are described in depth in the Implementation chapter below.

The emphasis in building this system was on simplicity – how simple could I make a mechanism that was nevertheless clearly recognizable as a social relationship? For example, the emotional model in the AlphaWolf system is limited to a single-axis model representing dominance; despite this sparse emotional information, the model succeeds in capturing the essential attributes of a social relationship. Because of the admitted simplicity of the model, the project was built with an eye to scalability along a variety of dimensions. Scalability will be addressed further in the Future Work section.

## **1.6.2 Directability**

The primary goal of this project is the creation of computational social relationships that resemble biological social relationships; nevertheless, there are certain inevitable differences between biological and computational relationships. Most central among these differences is directability [Blumberg 1995]. Directability entails some way for an external force to control the social relationship. In a social computer game character, for example, a player might wish to specify the relationships that are being communicated by that character. In multi-robot systems, too, as I will discuss in greater depth later, a person might wish to alter the social structures among the individual robots in a reliable way. As I mention above in the Research Problems section, directability of computational social relationships is the secondary focus of this research.

To make these social relationships directable, I drew a division between action and emotion; in AlphaWolf, participants are in control of the pup's action, but the pup autonomously maintains its emotional state. This distinction allows users to have high-level "intentional" control [Johnson 1999] over what the wolf does, but nevertheless maintains the wolf's independent personality by allowing it to have a continuous emotional attitude about the events in which the user is causing it to participate.

## ***1.7 Intellectual Framework***

This research project is situated at the juncture of a variety of disciplines. The social synthetic characters that I have built show influences from the fields of biology, artificial intelligence, traditional animation, computer graphics, and human-computer interaction. I discuss specific examples from these disciplines in the Related Work section of this document. In this section, I describe how this research (and in fact the research of the whole Synthetic Characters Group, to varying degrees) reflects each of these fields, and how it might have an impact on them.

As I mentioned above, the Synthetic Characters Group makes virtual creatures from natural models. The biological sciences are central to this process; we examine studies of real animals in order to ground our models in working real-world systems. Animals are functional systems; otherwise, they wouldn't continue to exist. The challenge of biology is to figure out how those systems work. Biologists produce models of the phenomena that they study; we use these models as the basis for our computational entities. If we are successful, our entities will provide synthetic support for the models that we've used.

The study of artificial intelligence is also central to our research. While building virtual canines may not seem directly relevant to building computer systems that are as smart as people, it is our belief that the kinds of abilities that dogs have are among the necessary underpinnings of human-level intelligence. There is a growing view among researchers of artificial intelligence and other disciplines that emotions and social competence are central to the kinds of behaviors that we call intelligent (e.g., [Byrne 1988; Damasio 1994; Dautenhahn 2000]). Research into social and emotional entities like the virtual wolves should provide additional support for this viewpoint.

The century-long tradition of animation also lies at the core of this research. Animators create compelling experiences for human audiences through moving pictures and sound (e.g., [Thomas 1981]). Usually (though not always), these experiences involve building the minds and bodies of virtual entities – characters – in order to tell stories. These characters, and the stories around them, are nearly always social and emotional in spirit. In most cases, animated characters are perceived through the linear media of film and television (though computer games are rapidly growing in importance, and embodied agents are on the horizon (e.g., [Cassell 1999])). Hopefully the research in this thesis will help enhance animators' ability to create characters and tell both linear and non-linear stories.

Our virtual entities find their shape through computer graphics – the art and science of creating beautiful images with digital systems. The graphical form that our characters take insures that our mechanisms "work as advertised." Because of their exposure to real animals and animated characters, people interacting with our wolves are able to determine whether or not their behaviors are explicable and appropriate. The graphical implementation of the AlphaWolf installation has allowed people to experience, and therefore evaluate, this research. The presence of the AlphaWolf installation at SIGGRAPH demonstrates that the work is already relevant to computer graphics.

The final discipline from which this research arises is the study of human-computer interaction (HCI). This field examines paradigms of interaction between people and computers, and seeks to design interfaces that take advantage of human abilities. The goal of this thesis has been to design a mechanism by which characters can form social relationships; if these characters are ever to form relationships with people, they need to be designed with an awareness of the ways by which those people might try to interact with them. While the primary relationships in AlphaWolf are those among the virtual wolves, the mechanisms described here could also form the basis for characters who interact with people face-to-face.

In approaching the problem of social relationships for synthetic entities, I have kept all of the above disciplines in mind.

## ***1.8 Scope***

The topic of synthetic social behavior, and even that of synthetic social relationships, could encompass quite a bit of territory; there must inevitably be some limitations on what is and is not covered by this research project. For example, it is well beyond the scope of this work to make full-size autonomous wolf robots who lurk around the Media Lab forming relationships with everything they see. In this section, I describe the limitations placed on this project.

The project is limited to wolves. Rather than attempting to make something that can compare to a human social relationship (an ill-fated venture at this point in our technological abilities), I've set my sights on the somewhat more accessible target of wolf social behavior. While some attempts will be made to draw comparisons to other social or potential social systems (e.g., humans, robots), the core focus of this work is the social behavior of the gray wolf.

This project involves only one small pack. Real wolves have both intra- and inter-pack interactions, and may have packs of as many as 29 individuals in size [Mech 1998]. Due to limitations in graphics rendering power, our pack never numbers greater than six, and never interacts with neighboring packs.

This project involves only a simplified subset of wolf social behavior. There is no breeding, no pup-rearing (except in terms of relationship-formation), no hunting, and only visual and auditory social behavior. (Although compelling, olfactory social behavior (e.g., scent marking) is beyond the scope.) The social relationship mechanism is vastly simplified compared to that of real wolves (though attempts are made to point the way towards greater complexity at various points in this document).

The project is a caricature of wolf social behavior. Real wolves, for example, spend much of their time asleep. To create a more interesting interaction for viewers, and to increase the rate at which they work out their relationships, our wolves sleep much less frequently. I have attempted to remain faithful to the essential spirit of wolf social behavior while modifying some of the details.

The project is limited in temporal scope. None of our packs has existed for more than a few hours of virtual or real time. While, in the AlphaWolf installation, this length of time may encompass a number of generations, there is no heredity among these generations.

Despite these limitations, a significant effort has been made to create an authentic, if stylized, replica of the natural model. In addition, I hope that some of the lessons learned will apply to a much wider range of synthetic social systems.



Figure 1-4: A gray wolf pup looks at its sleeping sibling.

## 1.9 Summary

The central hypothesis that lies at the heart of this research project is this:

- **The AlphaWolf mechanism is a novel, simple, robust, social relationship mechanism that can be used as the basis for a variety of relationships among machines and people.**

Throughout the rest of this document, I verify and support this central hypothesis through implementation details, quantitative evaluations, and thoughts on the usefulness of the AlphaWolf system in a variety of domains.

- Chapter 2 describes the AlphaWolf installation in depth, explaining the various elements that went into creating the piece shown at SIGGRAPH.
- Chapter 3 offers descriptions of related work and research projects that have formed the basis for this effort, and have informed many of the decisions made. This chapter situates the work in an intellectual context and confirms that it is original research. The two main areas addressed are biological systems (in particular wolf social relationships) and computational systems that pertain to social simulation.
- Chapter 4 gives details of the implementation of the social relationship mechanism and the three projects that have featured it. It provides overviews of the related sections of the Synthetic Characters code base. In addition, it describes the tools we use to visualize each applicable section of our code base.



- Chapter 5 describes the three ways in which I evaluated the success of the AlphaWolf system. First, I performed a set of human user studies to examine the effectiveness of the AlphaWolf at representing social relationships and its effect on users' subjective experiences. Second, I performed a set of simulations to explore the use of the AlphaWolf mechanism in the domain of synthetic resource exploitation. Finally, we observed participants at SIGGRAPH interacting with the AlphaWolf installation, and accumulated anecdotal evidence of their experiences.
- Chapter 6 discusses these key insights from these evaluations, and the limitations of the system.
- Chapter 7 offers a number of areas of future work, and a selection of possible real-world applications. The main areas discussed are entertainment, education, robotics, synthetic biology, fine art, computer-mediated communication technologies, human-computer interaction and artificial intelligence.
- Chapter 8 draws a range of conclusions from this research project as a whole.

Social relationships are an important part of human life. This dissertation presents an original mechanism by which computational entities may form social relationships, engage people in those relationships, and facilitate relationships between people. As I show in the coming chapters, the mechanism is novel, simple and robust, and could be used in a variety of applications. Although it is just a first step towards computational entities with the kind of social abilities found in people or other animals, perhaps it will help inspire future research on the topic. This project will have succeeded when it helps computational entities reap the benefits of sociality, to participate (albeit slightly at first) in the sphere of human social interactions, or to enhance the relationships among people.

## 2 The AlphaWolf Installation

The first major group of computational entities to be built with the system described in this thesis premiered under the title “AlphaWolf” (see Figure 2-1) in the Emerging Technologies program at SIGGRAPH 2001. The installation features a pack of directable 3D animated wolves. Three participants direct the actions of three pups – one gray, one black, and one white – in a newborn litter. By howling, growling, whining or barking into microphones (the “Howl Interface”), participants can tell their pups to howl, growl, whine or bark (see Figure 2-2). 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 2-1: The AlphaWolf logo.**

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



**Figure 2-2: A participant (Michael Wahrman) directs the actions of his wolf.**

Each participant stands at a different screen, on which a custom-built automatic cinematography system displays a view of his or her pup (see Figure 2-3). The three kiosks are networked together so that they can interact with each other. In addition to the three kiosks where the three participants interact, the full installation includes a fourth screen, a large (52”) plasma display, which displays a wider view of the pack’s interactions. This screen allowed a larger audience to view the wolves’ interactions.

## 2.1 Sample Interaction

To capture the feeling of the installation, I'll relate a hypothetical interaction episode in which a male participant directs the black pup.



**Figure 2-3: The full AlphaWolf installation, with three kiosks and a plasma display.**

At the beginning of the interaction, the black pup is asleep (see Figure 2-4) and the participant is standing in front of one of the three kiosks that make up the installation. The virtual camera shows a fairly close shot of the sleeping pup, and a quiet snoring comes out of the nearby speaker as the pup's chest rises and falls slightly. On the advice of an experienced bystander ("You have to wake him up..."), the participant picks up the microphone and says, "Wake up!" The pup's ear twitches at the sound. "He only really recognizes howls, growls, whines and barks," the guide says, so the participant gives a tentative bark into the microphone. The pup shifts in its sleep, its ear twitching each time it hears a noise into the microphone. After another noise, the pup opens its eyes and looks at the participant. The participant smiles, and howls. The pup stands up and howls too. The participant is starting to get the hang of the microphone interface.



**Figure 2-4: The pups start off asleep.**

The guide points out that he can also use the mouse to tell his pup where to go. The participant shifts the mike to his left hand, and clicks on the screen, a little way away from the pup. The pup turns and walks over to sniff the spot where the participant clicked. The user clicks a few more times, and the pup starts walking around the world to the various places he clicked.

As the pup walks around the virtual forest, another wolf pup, this one gray, becomes visible through the mist. The gray pup is just waking up. The participant hears a howl off to his left in the real world, and looks to see a woman at another kiosk with a microphone. The gray pup stands up and howls. “So she’s controlling the gray pup,” the participant realizes. Hearing the howl, the black pup turns its head to look at the gray pup. When it sees its sibling for the first time, an icon of the gray pup appears at the bottom of the black pup’s screen, and the black breaks into a gallop. The participant clicks on the gray pup to tell his pup to go visit.

As the black pup approaches, the gray pup turns to engage it. The black pup’s participant whines into his microphone to tell his pup to submit to the gray pup. When black reaches its sibling, it starts to perform a behavior called active submission among wild wolves, whining and pecking at the other pup’s mouth. The woman at the other kiosk growls into her microphone. Because of the woman’s growl, the gray pup starts to growl at black. Seeing gray’s dominant behavior, black switches to a passive-submit behavior, rolling over on its back, tail between legs. Before long, the woman decides to relent, and clicks her mouse in the distance to tell her pup to explore. The gray pup disengages, stops growling, and runs off dominantly (because of its uncontested domination of black). Black gets up submissively and, in response to its participant’s click, runs the other direction. The icon at the bottom of black’s screen has changed to show the gray pup standing in a dominant pose, reflecting black’s emotional memory of gray. Gray’s screen has a button for black, too, with black rolled over on its back in submission. The two buttons capture the two sides of the same relationship.



**Figure 2-5: Black dominates white.**

Over the next several minutes, the two pups meet their third sibling (who is white) and the adults in their pack. Each time they meet an individual, they form a relationship; when they encounter that individual again, they recall that relationship. In this hypothetical run, when black meets its white sibling, black’s participant gives a strong growl. Black, therefore, dominates white (see Figure 2-5).

The next time black encounters gray, even though it is just coming from a dominant interaction with white, it becomes more submissive simply because gray is present (that is, because black remembers being submissive to gray on a previous encounter). At the end of each interaction episode, the pups revise their relationship model to reflect the most recent interaction.

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. If directed by its user to growl at the gray pup, for example, the black pup will approach it submissively, but, despite the cringing attitude, nevertheless growl at it. The pup will perform the actions that its user directs, but in a style that demonstrates the pup's own impression of its relationships. As I show in section 6.1.4, the ability to direct a character to take an action that is "against its will" and yet have that character remain in character is one of the strengths of the AlphaWolf system.

The hypothetical interaction described above demonstrates the three different kinds of interactions that are formed during AlphaWolf. First, each participant builds a relationship with his or her pup. This relationship is usually collaborative, as each user works out how he or she wants the pup to interact with its pack mates, and the pup, through its relationships, tries to make that easier for the user. Second, the virtual wolves actively form relationships with each other throughout the run; these relationships are clearly displayed on the screens, and can be seen to change over time. The third relationship occurs between the various participants; for example, the man and the woman formed a kind of relationship on account of the interactions of their pups. In addition, participants and bystanders frequently started chatting because of the installation.

In the next few sections, I provide short descriptions of several main elements of the installation. For full implementation details, please see Chapter 4.

## ***2.2 Emotibuttons***

As mentioned above, each time a wolf pup meets a member of its pack, a button appears on the border of that user's screen. This button has an image of the other wolf on it. 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. If a user clicks on one of the buttons, the pup will run to the last place where the pup saw that wolf. 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.

One of the benefits of the emotional memory mechanism and the emotibuttons as a visualization tool is that they keep 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.

## ***2.3 Graphics***

The wolves are rendered using a custom "charcoal renderer" written by Marc Downie. 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. The unique graphical look of the AlphaWolf installation is one of its most immediately apparent strong points.

## 2.4 Sound

The background sound in the wolves' virtual world is of wind blowing through trees. In addition, each wolf makes a variety of sounds when it performs various actions. Walking is accompanied by the sound of footsteps; sleeping wolves snore; howling, growling and whining all make sounds.

The wolves' sounds change as they grow up, as well. Young wolves, for example, have high-pitched howls. Older wolves howl more deeply. While this change in the soundscape is subtle, it helps to convince participants that the wolves are really growing up.

While sounds clearly play a supporting role in this installation, it becomes quite clear how important they are when they're missing. The few times we have run the installation without sound, the people participating have seemed significantly less engrossed.



**Figure 2-6: The three pups, grown up now, have a final howl together.**

## 2.5 Development

During the five minutes of the installation, each pup grows up into an adult (see Figure 2-6). The pups grow up in terms of their relationships, finding their places in the social order of the pack. And they grow up physically, gradually taking on the shape and actions of an adult. When the interaction finally fades to black, the semi-autonomous pups, who had just been controlled by participants, become the autonomous adults for the next generation, and a new litter of three pups is born. The aging of the virtual wolves is based on the development cycle of natural wolves, although the exact timing is somewhat different; a pup's first five minutes correlate to the age of two to eight weeks in real wolf pups, and their second five minutes correlate to an arbitrary period of adulthood. For a more complete description of the wolf life cycle, please see section 3.1.1.1.

The development of the pups matched the learning curve of novice users. In the first few minutes of each run, participants are just figuring out the interface and getting their bearings in the virtual world. This period of adjustment is one of the main reasons that we caused the pups to start off asleep; having to wake up the pup without having to deal with the complexities of social

interactions gives people a simple introduction to the installation. Having the microphone be the only relevant interface while the pup is asleep also simplifies the introduction to the installation.

In the second half of each five minute run, people have usually mastered the interface and become more interested in the social phenomena that are going on among their pups. This shift in focus corresponds loosely to a transition that occurs in real wolf development, during which pups start to concern themselves with dominance interactions when they are several weeks old [Fox 1971]. These developmental changes are addressed more fully in the Related Work and Implementation chapters.

## ***2.6 Cinematography***

There are two main cinematic paradigms employed by the AlphaWolf installation, both based on previous work by the author [Tomlinson 2000b]. The first is a single-wolf perspective, which is used in each of the three kiosks where people interact. This paradigm attempts to show a specific pup (gray, black or white), the space ahead of the pup, and any other wolves with whom the pup is interacting. The goal of the three cameras with this paradigm is to help the participants interact with and direct their pups.

The second paradigm is used on the large plasma display, and is not associated with a particular participant or pup. This camera view looks at whichever group of pups is closest together (and therefore likely to interact). Because of the large size of the screen, the virtual camera tends to stay further from its target. This camera view allows a larger audience to get a general view of what is happening in the pack.

## ***2.7 A Networked Pack***

One of the important elements of AlphaWolf is that it allows three participants to interact at the same time. This multi-user aspect causes AlphaWolf to encourage the formation of relationships among the participants in the real world. Because of the cross-over between the virtual world and the real world, people can determine with whom their pups are interacting. It is not uncommon for participants to look around and say, “Who’s growling at me!?” or otherwise engage their co-interactors in a parallel interaction to the one occurring in the virtual world. It has become clear from watching people interact with the installation that part of the reason AlphaWolf was so well-received was this multi-player element.

## ***2.8 The Physical Set***

The physical layout of the installation also played a significant role in people’s experiences with AlphaWolf. The positioning of the kiosks with respect to each other affects the relationships that participants form with each other.

On the first day of SIGGRAPH, we had the installation arranged with the three kiosks fairly far from each other (~15 feet). After watching people interact with it as if it were three single-user installations, rather than a networked litter, we moved the three kiosks much closer to each other, so that they ended up essentially right next to each other (see Figure 2-3). This new configuration greatly enhanced the cross-over from the virtual world to the real world, bringing the simultaneous interactors into much more direct contact. People who were exploring SIGGRAPH together could readily play two pups in the same litter and have a chance to meet each other in this new format. Also, people who had never met, but whose pups formed relationships, would often end up chatting with each other in the real world after their run of the AlphaWolf installation was finished.

In addition, the traffic flow of the audience through the installation space had a certain effect on the way people interacted. If a person had taken a few moments to watch someone else interact, he or she tended to have a better interactive experience because he or she already knew the basics of how the experience would go.

## ***2.9 The Audience***

Over the course of the six day run of SIGGRAPH, between 500 and 1000 people interacted with the wolves, and another few hundred watched without participating (if, for example, there was a large crowd waiting to play). In addition, several hundred more people have interacted with AlphaWolf in its long-term installment in the Synthetic Characters laboratory at the Media Lab in Cambridge. These audiences have ranged in age from small children to senior citizens, have been of both sexes, and have come from many cultures and a variety of backgrounds. While children seem to take to it more rapidly than older people (as seems to be the case with many forms of technology), there has been a surprisingly consistent ability to “get the hang of” the interaction regardless of other factors. The interaction paradigm appears to cross cultural and language barriers with relative ease.

## ***2.10 The Video***

A short video (~3.5 minutes) of the AlphaWolf installation at SIGGRAPH can be found online at the following web address:

<http://web.media.mit.edu/~badger/alphaWolf/alphaWolf.mov>

It depicts a group of people interacting with the installation at SIGGRAPH, and shows a minute-long continuous clip of the wolves themselves.

## ***2.11 The Pitch***

In January 2001, seven months before AlphaWolf premiered in its final form, a preliminary sketch of the final installation was created for submission to the SIGGRAPH Emerging Technologies committee. While the wolves were much simpler behaviorally, the unique graphics renderer was already starting to take shape (see Figure 2-7), and the Howl Interface was functional.



**Figure 2-7: A pup tries to wake up an adult in the AlphaWolf pitch, January 2001.**



An interesting aside is how intelligent people *thought* the wolves were during the very early stages of AlphaWolf, when they were in fact quite simple. Creating *apparent* behavioral complexity can be quite easy (just have several wolves all walk to the same beacon that moves every few seconds and lie down when they get there.) Replacing this apparent complexity with *actual* behavioral complexity is the central goal of the Synthetic Characters Group.

## 2.12 The Team

AlphaWolf was created by a group composed mostly of MIT students over the period of a few months. In addition to the author of this document, the team included: Marc Downie, Matt Berlin, Jesse Gray, Adolph Wong, Robert Burke, Damian Isla, Yuri Ivanov, Michael Patrick Johnson, Derek Lyons, Jennie Cochran and Bryan Yong. Dan Stiehl, Rusmin Soetjipto and Dan Zaharopol helped create the initial version of the installation that was submitted to the Emerging Technologies committee in January 2001. Finally, no part of this project would have been possible without the continuing support, enthusiasm and guidance of Professor Bruce Blumberg.



Figure 2-8: Some Synthetic Characters make some synthetic characters.

### 3 Background and Related Work

The essence of the social relationship mechanism described in this thesis involves a number of different disciplines – biology, artificial intelligence, human-computer interaction, etc. Each of these topics has been well-studied by other researchers; together, they form the foundation on which the research described here is built.

The relevant background research can be divided loosely into two sections – natural systems and computational systems.

#### 3.1 *Natural Systems*

Natural systems do many of the things that we want computational systems to do. In order to make computational entities interact with each other, for example, it is useful to examine ways in which natural entities do it. In addition, as Dautenhahn points out, “[a]rtificial social agents (robotic or software) which are supposed to interact with humans are most successfully designed by imitating life.” [Dautenhahn 2000] Since living things are the complex systems with which we are most used to interacting, they provide a good model for making other complex systems that can engage us. Therefore, I begin with this section with a discussion of the natural social systems and other real-world phenomena on which this research is based.

##### 3.1.1 The Gray Wolf

This thesis attempts to create a mechanism for social relationship formation that is general enough and simple enough to be applicable to a wide variety of social systems; nevertheless, the specific implementation described here is based on one specific natural example – the gray wolf (*Canis lupus*).

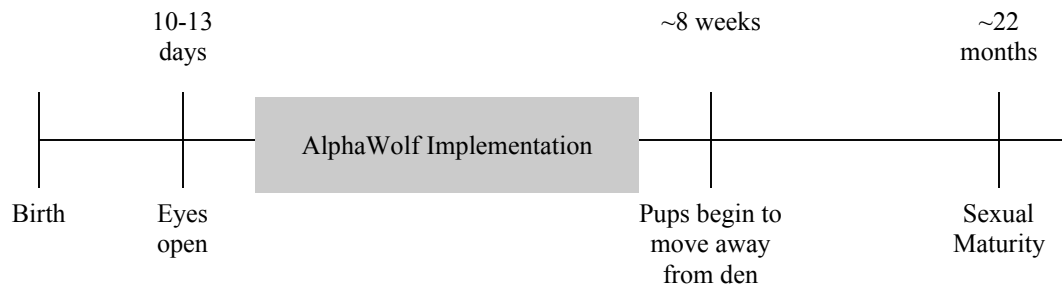


Figure 3-1: Two gray wolves, one gray in color, one black, having a dominance interaction.

##### 3.1.1.1 *Life Cycle*

Gray wolves are carnivorous mammals who live in parts of North America, Central America, Europe and Asia. They are closely related to the domestic dog, and are generally considered to be the evolutionary ancestor of dogs [Coppinger 1995]. Despite their name, gray wolves may

vary in color, with individuals being white, black, blond, gray, or combinations thereof (see Figure 3-1). Adult wolves weigh between 50 and 115 pounds on average, depending on sex, geographical location, and other factors [www.wolf.org 2000].

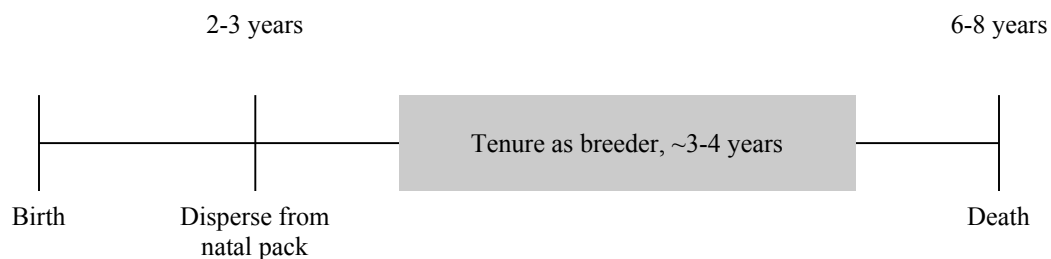


**Figure 3-2: Developmental stages of the gray wolf (not to scale).**

In their natural environment, gray wolves form complex social groups called packs. Packs of wolves live and hunt together. 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 1998]. 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.

Wolf pups are born in a den in mid-May after a gestation of approximately 63 days. Litters average approximately four to six pups in size (though litter size varies with prey availability and other factors). Pups are blind and deaf at birth, and nurse four or five times a day [Busch 1998]. Newborn pups only urinate and defecate when their mother licks their belly, a behavior known as reflex urination. By consuming the excrement, the mother keeps the den clean [Fox 1971].

Pups' eyes open at 10-13 days, and by three weeks they can hear and are able to explore the den area. Pups spend their first eight weeks in or near the den [Busch 1998]. Throughout this time, they are in continuous close contact with their litter mates. (See Figure 3-2.) The implementation described in this thesis focuses on this period of development. The rationale for this focus is addressed further in the Social Relationships section below.



**Figure 3-3: Timeline of an average gray wolf's life (not to scale).**

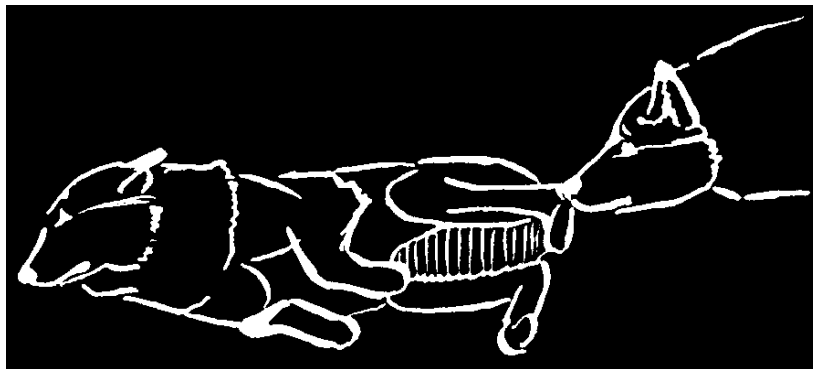
When the pack has young pups, the adult wolves travel away from the den to hunt, and carry back meat in their stomachs to feed the pups. Upon their return, the pups perform a stereotypical food-begging behavior, in which they crouch in front of an adult and lick or peck at the adult's muzzle.

This pup behavior incites the adult to regurgitate the meat, which the pups excitedly consume [Schenkel 1967]. Sometimes, shortly after regurgitating, the adult will re-consume the meat. This has the effect of training pups to grab hunks of meat and run off with them, rather than eating slowly at the site of regurgitation [Mech 1988].

Most young wolves disperse from their natal pack in their second or third year to find a mate and begin their own pack [Mech 1998]. The tenure of wolves as breeders varies from one to eight years, and averages three to four years [Packard in press]. Wolves live an average of six to eight years, and as long as 13 years in the wild and 16 years in captivity [www.wolf.org 2000]. (See Figure 3-3.)

### ***3.1.1.2 Social Relationships***

Within the pack social structure, individual wolves form long-term dyadic relationships with each other. These relationships are reflected in the behaviors of the wolves towards each other. Wolves' relationships change as a result of the interaction history between the dyad, and as such are an ideal model for the computational framework presented in this thesis. The ever-changing nature of the social organization of the pack provides a compelling evolutionary reason for wolves to require some cognitive mechanism for learning and remembering social relationships.




**Figure 3-4: Passive submission (from [Schenkel 1967]).**

On a developmental time scale, social relationships begin to form as pups' senses come on line (i.e., approximately 13 days after birth). Wolf social development is differently timed than dog social development, in that dogs' "window of social opportunity" (i.e., the period in which they decide what kinds of entities "count" as social partners) may last until they are 16 weeks old, while that of wolves is rapidly closing by 19 days of age [Coppinger 2001].

Since the focus of this thesis is on how social relationships come about, the most relevant section of the wolf life cycle is the period when most social relationships are first formed. Therefore, the "lead characters" in the virtual packs of wolves will be pups between 19 days and eight weeks of age. (See Figure 3-2.) During this period, they are old enough to have fully formed sensory apparatus and to be able to identify wolves as social partners, but young enough that they are just working out their relationships with their pack mates (and in particular with their litter mates).

Wolves exhibit a wide range of behaviors to communicate their relationships with social partners. Dominance and submission behaviors are central to these relationships. There are two main types of submission that wolves exhibit – passive submission (see Figure 3-4) and active submission (see Figure 3-5). Passive submission involves a wolf lying on its side or back, exposing the ventral side of its chest. The ears are held close to the head, and the tail is tucked between the legs. Less severe forms of passive submission may involve looking away and other non-

threatening behaviors. Passive submission usually occurs in response to dominance behavior (see below) by a social partner whom the submitting wolf believes is dominant to him [Schenkel 1967].



**Figure 3-5: Active submission (from [Schenkel 1967]).**

Active submission involves a crouched posture with backward directed ears, and licking or pecking the mouth of the dominant wolf. This set of behaviors usually occurs in interactions initiated by the actor toward a more dominant social partner. It is often matched with tolerant behavior on the part of the dominant individual [Fox 1971].

Dominance behavior involves a “confident” posture with tail and ears erect, direct staring at the social partner (see Figure 3-4). More vigorous dominance behaviors involve pinning the muzzle of the submissive wolf to the ground and growling [Schenkel 1967].

In response to active submission, a dominant individual usually offers tolerant behavior. “If active submission is displayed, parent may respond to cub, or superior adult to subordinate, by standing erect and aloof, accepting the display or greeting, and perhaps briefly lowering its ears, wagging its tail and ‘grinning’ slightly” [Fox 1971, p. 92]. (See Figure 3-5.)

Wolf social behaviors appear to be derived from other behavioral patterns exhibited by wolves [Schenkel 1967]. For example, passive submission bears a resemblance to the infantile behaviors involved in reflex urination, mentioned above [Fox 1971]. Active submission, also, is very similar to the food-begging behavior of pups, also mentioned above. (See Figure 3-6.) Similarly, dominant behaviors appear to be a form of “ritualized fighting” [Golani 1983].

### ***3.1.1.3 Research Efforts***

Researchers have undertaken a variety of long-term projects to study wolves. These studies fall roughly into two categories – field studies and captive research. Three researchers in particular are central to the ideas in this thesis.

L. David Mech has conducted the most thorough studies of wild wolves to date, e.g. [Mech 1998; Mech 1999]. He leads three significant studies of wolves in their natural ecosystems – a study of

wolf population trends and wolf-deer coactions in the Superior National Forest, Minnesota, begun in 1968, a study of wolf-population trends and wolf-caribou coactions in Denali National Park, Alaska, begun in 1986, and a study of wolf social behavior on Ellesmere Island, Canada, begun in 1986. Dr. Mech offers the most scientifically rigorous analyses of wild wolves available.

Michael W. Fox's studies on hand-reared and zoo-reared canids, in particular wolves, offer some of the most explicit descriptions of young wolf behavior that I have found. While some of his interpretation may not be consistent with current beliefs about wolf behavior, he is careful to separate observations from interpretations. Therefore, I rely heavily on his descriptions of the behaviors and interactions of wolf pups.



**Figure 3-6: Active submission in adults resembles food soliciting behavior in pups.**

Erich Klinghammer, head of the Wolf Park captive wolf facility in Battle Ground, Indiana, has had extensive experience with rearing wolves in captivity. As part of the research for this thesis, Bruce Blumberg and I attended his six-day seminar on wolf and dog behavior and training (see Figure 3-7); Klinghammer's thoughts on wolf socialization are prevalent throughout this work.

A relevant distinction to bring up is the behavior of wolves in captivity versus the behavior of wild wolves.<sup>1</sup> In the wild, as I mentioned above, the wolf pack is essentially a family, with one breeding pair, their pups and yearling offspring. In captivity, on the other hand, multiple adult wolves are often kept in the same enclosure, thereby causing more social unrest than might be seen in the wild (where wolves may disperse rather than fight). As David Mech points out, "in natural wolf packs, the alpha male or female are merely the breeding animals, the parents of the pack, and dominance contests with other wolves are rare, if they exist at all. During my 13 summers observing the Ellesmere Island pack, I saw none." [Mech 1999, p. 4] Despite this distinction, both kinds of interaction reflect actual behavior of biological wolves. Because of the lack of dispersal mechanism, the virtual wolves described here exhibit behaviors closer to those exhibited by captive wolves, who must negotiate relationships with a specific set of wolves, without the option of dispersal.

---

<sup>1</sup> Wolves are certainly not the only species to exhibit different social behaviors in captivity and in the wild. Chimpanzee social life, too, for example, is altered by captivity. (de Waal 1994)





**Figure 3-7: The author with a captive gray wolf at Wolf Park in Battle Ground, Indiana.**

### **3.1.2 Social Relationships**

While the focus of the research in this thesis is on creating a system that can mirror the kind of long-term social structures found in wolves, it is important to have a broader view of the range of social relationships that occur in natural systems. For example, making an installation that can interface with human participants requires some consideration of how humans and other primates form relationships. In addition, to create a system that is not so tailored to one function that it is unable to serve any other, it is relevant to keep in mind other kinds of relationships that we might want to implement in the future.

Many animals actively change their environment, or their relationship with other individuals in their environment, to make that environment more predictable. Cohen, Riolo and Axelrod [Cohen 1999] refer to consistency for an individual within an environment as the *preservation of context*. Many actions that organisms take can be seen as mechanisms for context preservation. Social structures, in particular long-term dyadic social relationships such as those formed by wolves and primates, serve to preserve the context in which individuals must function, thereby enhancing their ability to predict and exploit their environments.

Eusocial insects are masters of preserving context; they manipulate their physical environment in ways that cause their environment to be more stable (termite (e.g., *Macrotermes bellicosus*) nests are engineering wonders with respect to thermoregulation [Wilson 1980]). In addition, the very similarity of all the members of each caste serves to greatly limit the range of interactions that a given termite will encounter. Thus, both the physical and the social contexts within a termite nest are quite consistent.

Group living, as among wolves (see above) or chimpanzee communities (*Pan troglodytes*) [Wrangham 1994] provides for certain elements of consistency over time. Many evolutionarily beneficial phenomena can occur without the necessity for individuals to know each other, such as dominance hierarchies based on direct perception rather than persistent memory [Hemelrijk 1996; Hemelrijk 2000], kin selection [Maynard Smith 1964] and reciprocal altruism [Trivers 1971; Packer 1977; de Waal 1992]. (More recent studies, (e.g., [Mitani 2000]) suggest that kin selection may not be as important as was once thought, at least with respect to chimpanzee social groups, but still allow that it plays an “ancillary role” ([Mitani 2000], p.885).) Although in practice it is difficult to disentangle the effects of proximity and genetic relatedness from the

effects of remembered relationships, it is possible for these phenomena to occur in the absence of remembered relationships. Many characteristics that seem to be the result of long-term dyadic relationships may be emergent phenomena from interactions of simpler elements (e.g., male group formation in simulated chimpanzees [teBoekhorst 1994], mobbing behavior in spotted hyenas (*Crocuta crocuta*) [Zabel 1992]).

Just as context could be preserved by means of relatively static neighborhoods in which neighbors interact repeatedly, long-term social relationships allow each individual to maintain her own *social* neighborhood. It is the persistence of interactions, not physical proximity, that is the crucial element.

As mentioned in the Introduction, long-term dyadic social relationships enable the existence of stable alliances that we see in primate societies [Harcourt 1992; Barrett 2000]. Social relationships make possible the raiding parties found in chimpanzees [Wrangham 1996]. Pair-bonding, among wolves for example [Fox 1971], relies upon a persistent relationship between two individuals, and in the case of wolves provides for at least one other individual to hunt while the mother is in the den with newborn pups. Long-term supportive relationships also occur between mothers and offspring in such species as spotted hyenas and cercopithecine primates [Engh 2000]. These relationships appear to allow offspring to inherit their mothers' ranks by means of maternal intervention. Kawai [Kawai 1958] makes a distinction between basic rank, which reflects the relationship between two individuals on their own, and dependent rank, which is how those individuals would fare within their social context (in Kawai's case, how the Japanese monkeys compete in the presence of their mothers.) Long-term dyadic relationships cause dependent rank to be more important than basic rank in determining access to resources.

In a population that exhibits long-term dyadic relationships, additional evolutionary pressure is put on communication. Greetings, negotiation and recruiting all require a vocabulary of signals [Hauser 1998]. Humans are masters of interpreting social signals and meaning [Brothers 1997]. The ability to sense emotional states takes on greater subtlety. No more do the same basic signals capture a full dynamic range of communication. The longer a chain of interactions will be sustained between two individuals, the more complex their communication system will need to be. In addition, organisms who participate in these relationships need to spend time and energy maintaining the social bonds (for example by exchanging a "social currency" such as grooming [Barrett 2000]), so that the relationship will be firmly in place when it is needed.

Among animals who tend to live in groups, building relationships with conspecifics conveys a significant evolutionary advantage. Relationships are reciprocal patterns of reliability. If I have a relationship with someone, I am more able to predict what he is likely to do, and to alter my strategy accordingly. By archiving my experience with him (and he with me), and acting on that stored information in the future, we are able to preserve the context of our interaction, rather than starting each new meeting afresh. This carries with it some evolutionary costs (e.g., my ability to communicate deceptively is somewhat curtailed). The benefits, however, of social maneuvering by means of long-term relationships have caused it to persist in a wide variety of mammalian social systems.

Social phenomena do not necessarily increase the fitness of the population as a whole, with respect to its environment. Rather, they serve to benefit certain individuals (those skilled at relationship-building) at the expense of other individuals who might have prospered under other circumstances (e.g., the biggest and strongest). Skill at grooming [Barrett 2000] begins to compete with skill at fighting as a means of resource-access. The advent of long term dyads changes the social landscape on which all the members of that society exist. As soon as a few



individuals form alliances, even the strongest individuals will also need to adopt an alliance-forming strategy in order to maintain their position.

### 3.1.3 Emotion

The mechanism of social relationship formation presented in this thesis is inherently emotional in nature. Therefore, an understanding of the scientific study of emotion is relevant. Charles Darwin's ideas about emotions [Darwin 1965 (originally published 1872)] form the basis for much of modern research into understanding emotions. He focused on the *expression* of emotion – the ways in which emotions manifest themselves and are communicated among individuals.

For the AlphaWolf project, we considered two main emotional paradigms – a categorical approach and a dimensional approach. The categorical approach separates emotional phenomena into a set of basic emotions – for example, fear, anger, sadness, happiness, disgust and surprise [Ekman 1992]. While this approach captures the array of human emotions, it is more extensive than we needed for the behavior of the wolves. In addition, since there appears to be significant overlap among the categories (e.g., happiness and sadness lie at opposite ends of a similar phenomenon), this model doesn't lend itself as well to the kind of computational implementation that we wanted for our wolves.

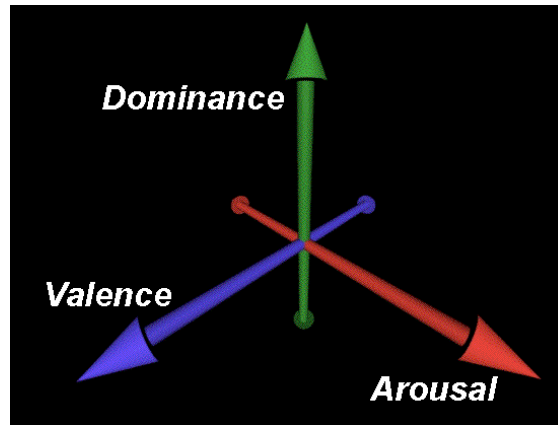


Figure 3-8: A dimensional emotional model.

The dimensional approach (e.g., [Schlosberg 1954; Mehrabian 1974; Smith 1989; Plutchik 1991; Russell 1997]) maps a range of emotional phenomena onto an explicitly dimensioned space (see Figure 3-8). Various dimensional spaces have been proposed; for example, Russell offers Stance, Valence and Arousal as axes describing a comprehensive emotional space [Russell 1980]. Mehrabian and Russell offer a similar model, with Pleasure, Arousal and Dominance as the canonical axes [Mehrabian 1974]. We found that this last model most effectively reflected the kinds of phenomena we have found in wolf social behavior. For example, Dominance is central to the relationships that pack mates form, Pleasure serves as the reward function by which the wolves learn, and Arousal accounts for the excitement level that wolves exhibit as their expectations of a situation change.

Computational implementations of both categorical and dimensional models will be discussed below in the “Computational Systems” section of this chapter.

### 3.1.4 Perception and Communication

The way in which social relationships are signaled depends largely on the perceptual abilities of the social partners. In order to understand how social signals work in a specific natural example,

it is relevant to attend to the perceptual mechanism of that species [Hauser 1998]. The four major senses that have a significant impact on wolf social behavior are smell, hearing, sight and touch.



**Figure 3-9: A wolf communicates.**

Wolves have a wide array of vocalizations, including “whimpering, wuffing, snarling, squealing and howling” [Zimen 1981, p. 68]. (See Figure 3-9.) 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. For example, scent is used by wolves and dogs to determine the unique identity of social partners [Beaver 1999], and for scent marking objects in the environment. The sense of touch plays a part in wolf social contact, in particular through biting, pinning, and other physical ways of showing dominance. In wolves, as in most social creatures, communication is central to the social relationships that are formed, and perception lies at the heart of communication.

Human perception shows certain weaknesses when compared to wolf perception. Humans are one hundred times less effective at smelling than wolves are [Busch 1998]. Wolves also have a more acute sense of hearing than humans. Wolf and human perception is much closer in the visual realm, however. Naturalist R.D. Lawrence believes that wolves are unable to distinguish pack mates beyond a distance of 100 to 150 feet [Busch 1998]. The goal of this thesis has been to make a biologically accurate computational implementation of wolf social behavior that is also intelligible to people and feasible to create on current computer hardware; therefore, it has been necessary to keep human perceptual abilities in mind.

### **3.1.5 Learning**

Evolution may not be able to provide mechanisms to cope with certain aspects of an animal’s environment that change rapidly or are otherwise unpredictable on evolutionary time scales [Plotkin 1994]. As Blumberg puts it, “The best that evolution can do in these situations is to provide mechanisms that make it easier for the animal to learn these kinds of predictable regularities” [Blumberg 2001 (to appear), p.4].

Social relationships inherently entail an element of learning, to allow individuals to remember each other from interaction to interaction. Here I mention examples of animal and human learning that are relevant to this thesis project.

### 3.1.5.1 Emotional Learning

The emotional memory mechanism presented in this thesis is based on Damasio's "Somatic Marker Hypothesis" [Damasio 1994]. This hypothesis proposes that people (and animals) attach emotional significance to stimuli that they encounter in their environment, and then re-experience those emotions when they encounter those stimuli on future occasions. Although the social relationship mechanism that I describe in the Approach section treats *individuals* as emotional significant stimuli, a stimulus does not have to be an individual – only a *causative entity* [Damasio 1994, p. 162]. This entity could be some conjunction of phenomena (e.g., two wolves at the same time, or a wolf who smells of meat).

Various other researchers have addressed the significant role that emotions play in animal and human learning [Buck 1984; Frigaszy 1996; LeDoux 1996].

### 3.1.5.2 Social Learning

Social learning is the process by which the "acquisition of behavior by one animal can be influenced by social interaction with others of its species." [Galef 1996, p.8] Social learning, as defined above, occurs in a wide variety of animal species. There are two loose categories of social learning – one involving social learning of behaviors that are not social in nature, and the other involving social learning of social behaviors.

Various species of animals learn non-social behaviors from their interactions with social partners. Social interactions factor into the acquisition of feeding habits in rats. For example, black rats (*Rattus rattus*) learn to pull the scales off pine cones by means of social transmission (i.e., being exposed to the partially stripped cones being eaten by their mothers) [Terkel 1996]. Galef, too, discusses mechanisms by which rats (in this case Norway rats (*Rattus norvegicus*)) may extract information about feeding from more experienced social partners, for example, by following individuals who smell like food [Galef 1996]. Capuchin monkeys are more likely to eat when there are social partners in the area, and are more willing to accept novel foods in a social context regardless of what the social partners are eating [Visalberghi 2000].

The learning of social behaviors, too, is influenced by social interactions. Bird song learning depends on social interactions, as Irene Pepperberg's studies with model-rival training show [Pepperberg 1999]. Social play lets wolf pups learn group dynamics that may contribute to their ability to hunt as packs [Bekoff 1974]. Chimpanzees learn which behaviors are socially appropriate by means of rewards and punishments delivered by their group mates [Goodall 1986]. All of these examples show how social interactions affect an individual's acquisition of social behavior.

One element of learning that is often studied is the topic of imitation. Humans, for example, are able to perform imitative learning, mirroring the behaviors of a social partner with their own behaviors. Unlike humans, wolves do not appear to perform imitative learning. Therefore, it is necessary that the model presented in this thesis does not rely on individuals being able to imitate each other. Thomas Zentall provides an excellent summary of nonimitative social learning [Zentall 1996]. He discusses contagion, social facilitation, socially mediated aversive conditioning, local enhancement, stimulus enhancement, observational conditioning, and following behavior, each a different kind of nonimitative social learning.

An interesting distinction that has some bearing on the research presented here is the one between horizontal and vertical transmission [Laland 1996]. Horizontal transmission is the acquisition of behavior by means of interactions among members of the same generation (e.g., among siblings).

Vertical transmission is the same phenomenon occurring across generations (e.g., from parent to offspring). Erich Klinghammer and Patricia Goodmann [Klinghammer 1985] describe that hand-raising a single wolf pup from around day 10 to the age of four months without exposing it to regular contact with conspecifics caused that pup to be unable to re-integrate into the social order of the pack. On the other hand, they have repeatedly raised two or more pups with each other, but otherwise isolated from lupine contact (i.e., no adults), and those pups have little or no problem being accepted back into the pack. This difference suggests that some kind of social learning or development goes on among young wolves that does not require experienced adult individuals (i.e., horizontal transmission). Coppinger describes that, during the “critical period for social development,” dogs “learn what species they belong to.” [Coppinger 2001, p. 106] As mentioned above, this period occurs in wolves from day 13 to day 19, within the window when Klinghammer and Goodmann raise the pups away from adults. If wolves are raised without other wolves during this critical period, they will not see other wolves as social partners, preferring humans instead. Allowing them to live with other puppies during this period gives them the necessary cognitive apparatus to interact successfully with adults throughout the rest of their life.

One element of human social behavior that is appropriate to the design of the interface to the AlphaWolf installation is the topic of social referencing [Siegel 1999]. Humans, especially children, in novel situations often look to social partners for cues about how to react. People interacting with the virtual wolf pups often drew their emotional cues from the pups themselves. This will be discussed further in the Evaluation chapter of this thesis.

### **3.1.5.3 Training**

Directability of a virtual entity’s social relationships is one of the central themes of this thesis; dog training is a clear example of directability in the natural world. By delivering rewards and other reinforcements, dogs can be trained to perform a wide variety of behaviors on cue [Serpell 1995; Coppinger 2001]. While the training paradigm is a bit different from the control paradigm used in AlphaWolf (which is more similar to electrodes into a virtual brain), it is nevertheless an interesting topic for comparison in terms of directing the behavior of an autonomous entity.

### **3.1.6 Development**

Development ties in to this research as well, since social learning is a large part of the developmental process. Evolution, development and learning can be seen as similar adaptive structures operating at different time scales [Dickins 2001].

Critical periods are times in development when organisms are more prone to developing certain abilities. The idea of critical periods (or the less sharply defined version, “sensitive periods”) of development for acquiring certain skills crosses a range of species boundaries. In wolves, a wide array of behaviors “come online” at around 22 months, when most individuals reach sexual maturity [Mech 1998]. Lorenz’s studies on imprinting in greylag geese reflect a critical period for social attachment [Lorenz 1952]. Matsuzawa proposed a critical period in chimpanzees for learning the skill of cracking nuts with rocks [Matsuzawa 1994]. Cats, monkeys and humans exhibit critical periods for visual ability [Gould 1982; Dowling 1992]. Dogs have critical periods for learning and practicing submissive behaviors [Coppinger 2001].

Predatory behavior in cats also exhibits a strong developmental element [Lorenz 1973]. Rather than cats having a single, unitary “prey-catching drive”, each of the subsets of prey-catching – lying in wait, stalking, chasing, seizing, killing, eating – has an “action-specific energy” [Lorenz 1973, p.223]. This action-specific energy lets the cat learn how to do each sub-behavior for its own sake. After the cat has learned the sub-behaviors, it can chain them together into the full

prey-catching behavior. This breaking up of one adult behavior into multiple sub-behaviors, each of which can be learned independently over a developmental time scale, allows the learning of an extremely complex behavioral pattern to become tenable.

## ***3.2 Computational Systems***

While biological systems provide the inspiration and direction for this research project, other computational systems that have been built are central to the method of implementation, evaluation, and application of this project.

### ***3.2.1 Synthetic Characters***

The previous research of the Synthetic Characters Group at the MIT Media Lab is the substrate on which this thesis has been built. Both the conceptual framework and the extensive code base that the Characters Group has created since 1997 are crucial to the social relationship mechanism presented here. Our research springs from an academic tradition grounded in the work of Minsky [Minsky 1986], Brooks [Brooks 1990] and Maes [Maes 1989].

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 [Burke 2001; Downie 2001a; Isla 2001b; Blumberg 2001 (to appear)]. I will mention a few of the systems here that are most central to the process of adding social competence.

#### ***3.2.1.1 Action Selection***

One of the most basic elements of our system is the action selection system. We have used several mechanisms for action selection over the years; the basic mechanism that we currently use, and that I use in this thesis project, is the “ActionTuple.” Each ActionTuple consists of four main 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. The action selection mechanism and other parts of our system are discussed more fully in Chapter 4.

#### ***3.2.1.2 Expression***

When an ActionTuple becomes active, it causes the motor system [Downie 2001a] to execute the Action itself. Each Action is executed in an emotional style; our system provides this expressive range by blending between example animations. This system is based on Rose’s “verbs and adverbs” system [Rose 1999], in which an action (a “verb”) is taken in a certain style (an “adverb”). Because our motor control system can blend between example animations, our animators need only create a few extreme emotional styles of each Action (for example, one dominant walk and one submissive walk) to get the full dynamic expressive range between those examples.

#### ***3.2.1.3 Learning***

Our system already performs several kinds of learning. The three main types of learning currently in place are state-space discovery, action-space discovery, and state-action-space discovery [Blumberg 2001 (to appear)]. State-space discovery involves a creature learning new configurations of its perceptual world that consistently lead to reward. Action-space discovery

entails learning new body configurations and actions that correlate with a reward. Finally, in state-action-space discovery, the creature connects the two, learning that certain combinations of a state-space and an action-space lead to a reward. All of these bias the action selection of the creature in accordance with Thorndike's Law of Effect [Thorndike 1911], which predicts that phenomena that correlate with positive results will be chosen more frequently in the future.



Figure 3-10: The raccoon in *Swamped!* at SIGGRAPH '98.

#### 3.2.1.4 Installations

*Swamped!* [Blumberg 1998] featured a virtual chicken running around a barnyard scenario on a large projection screen. Participants could control the chicken by means of a sensed plush toy (a bright yellow fleece chicken). A fully autonomous raccoon (see Figure 3-10) marauded around the barn yard in search of the chicken's eggs. The *Swamped!* installation featured an action selection mechanism, a motor system and a novel interface for interacting with autonomous and semi-autonomous characters [Johnson 1999].



Figure 3-11: Elliott the Salesman from (void\*).

(void\*): A Cast of Characters [Blumberg 1999] showed three humanoid characters sitting at an all-night diner. The interface, two dinner rolls with forks stuck in them (in the spirit of Charlie Chaplin's film "The Gold Rush"), allowed a participant to make the characters get up and dance. The characters were different from puppets, though, having emotional responses to the interaction that they were undergoing and changing the entire style of their animation and interaction to reflect their emotional state. (See Figure 3-11.) In addition, the characters in (void\*) could learn the ways in which people interacted with them, and would continue to act in those ways after the participant stopped interacting [Yoon 2000b].



**Figure 3-12: Duncan herds ornery sheep in sheep|dog.**

Finally, two of our latest projects feature an adaptive autonomous animated terrier named Duncan. *sheep|dog: Trial by Eire* showed at the opening of the MediaLabEurope in Dublin, Ireland (July 2000), and at the Electronic Entertainment Expo in Los Angeles (May 2001). This installation allowed a participant to play the role of a shepherd in a virtual sheep herding competition. (See Figure 3-12.) By means of his trusty terrier, the participant was able to coax a flock of ornery autonomous sheep around a field and into a pen. In *Clicker By Eire*, the participant, using dog training techniques borrowed from the real world, can train *Duncan* to perform a variety of tricks in response to verbal cues. Technically, *Duncan* is our platform for focusing further on learning, action selection and motor control in autonomous characters.

AlphaWolf is next in the succession of Synthetic Characters installations, extending our character-building capabilities to include social competence.

### **3.2.1.5 Other Systems**

Various other researchers have also addressed the topic of believable agents. Through both the Oz Project at Carnegie Mellon University and the company Zoesis, Joseph Bates and his colleagues have created computational characters who appear lifelike and are able to interact with people in real time. The focus of these groups is to build interactive storytelling systems. Most relevant is Scott Neil Reilly's doctoral research into believable social and emotional agents [Reilly 1996].

Ken Perlin and his colleagues in the Media Research Lab at NYU have done pioneering work in creating Synthetic Actors (e.g., [Perlin 1996]). By working closely with the natural style of their characters' motion, they have created virtual characters who move and interact very naturally.

Justine Cassell's Gesture & Narrative Language (GNL) research group at the MIT Media Lab builds virtual humanoids with the ability to express themselves like real people. Through various



projects including a virtual real-estate agent [Cassell 1999] and a virtual storytelling playmate [Ryokai in press], GNL has explored natural forms of communication for virtual entities.

The research by these three groups has focused primarily on ways of making the social and emotional interactions of virtual characters *believable* to people; the focus of this thesis is on making entities that form viable and long-lasting social relationships among computational entities *as well as* making characters who are socially engaging to humans.

### 3.2.2 Social Robotics

Perhaps the most compelling example of a social interaction between a human and a machine is the robot Kismet, created by Cynthia Breazeal [Breazeal 2000]. (See Figure 3-13.) This robot has the form of the head of a young fanciful creature with large features. Kismet is able to interact in real-time with people, engaging them in turn-taking conversations (with Kismet using baby-like vocalizations), and reacting appropriately to the interactions. The research underlying Kismet focuses on building computational systems that can take advantage of the natural set of behaviors that people offer to infants. These behaviors act as the scaffolding to help those infants learn. Perhaps by giving a computational system the benefit of this scaffolding, the problem of learning how to interact with the real world will become more tenable. This project will be discussed further in several sections below.



**Figure 3-13: Kismet.**

For a general introduction to social interactions among groups of robots, the reader is directed to Ronald Arkin's book *Behavior-Based Robotics* [Arkin 1998].

### 3.2.3 Game-Theoretic Models

As I discussed in the section on natural social relationships, social species create environments in which context may be preserved by developing social relationships that persist over long periods of time. Two individuals who know each other (i.e., can reliably distinguish each other from anyone else) can pick up their interactions where they left off, rather than starting afresh each time they meet. In a game theoretic model where the Prisoner's Dilemma is played only once,



the dominant strategy is to defect. However, in the iterated version, strategies such as Tit For Tat perform much better [Axelrod 1984]. Two individuals who know each other are playing an iterated game.

As I mentioned above, context preservation in a multi-agent system has been explored by Cohen, Riolo, and Axelrod [Cohen 1999]. In their Iterated Prisoner's Dilemma experiments, they found that the preservation of context enabled the emergence of social organization and cooperation. The simplest form of preservation of context (in a simulation) would be to cause each agent to interact with the same agent on every turn. For agents embedded in a two-dimensional space, preserving context entails "neighborhoods" in which there is some predictability with regard to whom each agent tends to encounter in successive turns. Cohen *et al.* found that, in fact, preservation of context can be more general than simple geographic neighborhoods, extending to mixing strategies as well as long as context (i.e. the social neighborhood) was preserved.

One possible evolutionary mechanism by which individual recognition could arise is by increasing the costs of deceptive communication. Deceptive signaling can occur in situations where the benefit to the deceiver outweighs the cost, as shown by Adams & Mesterton-Gibbons [Adams 1995] in a game theoretic model of the behavior of stomatopod crustaceans (*Gonodactylus bredini*). The ability to remember specific individuals makes it possible to bias one's behavior against them in the future if they are ever caught in a lie. This acts to reduce deceptive communication, by increasing the cost associated with deception.

### 3.2.4 Multi-Agent and Multi-Robot Systems

This thesis is far from the first time that a researcher has explored the notion of computational social behavior. While this section does not provide a comprehensive review of the literature, it does describe the most relevant projects pertaining to the research described in this thesis.

Craig Reynolds' Boids research [Reynolds 1987] has been one of the great inspirations throughout this thesis. In his 1987 SIGGRAPH paper, he presents an exceedingly simple model by which bird flocking could be simulated. His "boids" follow three simple rules (in decreasing order of precedence) – they avoid collisions with their flock mates, they match velocity with their flock mates, and they try to stay close to their flock mates. (See Figure 13.) From these rules, repeated in many individuals, bird flocking emerges. The purpose of this thesis is to arrive at a similarly simple mechanism by which social relationships can be formed among computational entities.




Figure 3-14: Craig Reynolds' Boids (from <http://www.red3d.com/cwr/boids/>).

Researchers in the USC Robotics Research Lab have done research in multi-robot coordination that is exceedingly relevant to the mechanism presented in this thesis. In particular, one research project [Vaughan 2000] presents a system that resolves interference among simulated multi-robot systems by means of a variety of aggressive mechanisms, with the losing robot backing away to allow the winner to pass. They compare three experimental mechanisms to a control group – a random winner of each encounter, a fixed dominance hierarchy, and a dynamic system based on the amount of space behind each robot. They found that, while all three aggressive mechanisms outperformed the control group, no one of the three was significantly superior to the others. “We hypothesize that a dominance hierarchy will be effective only when there are (i) non-uniform abilities in the group and (ii) a relatively slow change in the abilities of individuals. These conditions are relatively unusual in robotics with the important exception of systems with learning, evolution or other long-term adaptation.” [Vaughan 2000, p. 9] They also note that “with increasing population, task complexity and/or increasing uncertainty in the environment, conventional planning becomes less tractable.” [Vaughan 2000, p. 9] In the Evaluation section of this document, the mechanisms described in [Vaughan 2000] will serve as alternative approaches, to which I will compare the mechanism presented in this thesis.

The topic of multi-agent and multi-robot systems has become a significant subset of the research being done in autonomous agents (e.g., International Conference on Multi Agent Systems [Lesser 1995]). Dautenhahn [Dautenhahn 1998], for example, provides an analysis of human social intelligence for the purposes of making socially intelligent computational agents. Robot soccer competitions are another venue that requires the coordination of multiple entities [Stone 2000]. Many others have also done relevant research into the interactions of multiple autonomous entities (e.g., [Resnick 1994; Di Paolo 1999; Noble 1999], Bonabeau and others at the Santa Fe Institute [Cohen 1999; Bonabeau 2000], etc.).

Various other researchers have studied synthetic social systems from natural models, including fish schools [Tu 1994], gorillas [Allison 1996], chimpanzees [teBoekhorst 1994] and other primates [Bond 2000]. Our research differs from these efforts in our focus on emotion and learning as key components of social competence.

Hemelrijk [Hemelrijk 1996] did experiments using a similar social relationship mechanism to the one we use in the AlphaWolves, except that the memory mechanism did not have a confidence component, was not applied in a continuous fashion, and was not connected to graphical expression. Hemelrijk’s implementation was based on another by Hogeweg, whose implementation featured “Estimator” agents that tried to establish whether they might win or lose potential dominance interaction [Hogeweg 1988]. In addition to the differences listed above, AlphaWolf’s 3D animated visualization of the social computational entities makes the project significantly different from these works.

### **3.2.5 Perception**

Perception and communication are central to producing and understanding social signals. Various researchers have simulated the ways in which animals engage in social communication. Noble, for example, created evolutionary simulations of food- and alarm-calling, and of aggressive social signals [Noble 1998].

While perception is a central component of the model presented in this thesis, the inner workings of the perceptual mechanism is outside the scope of this research. The social relationship mechanism uses much the same perceptual mechanism that is described in [Burke 2001] and [Isla

2001a]. This perceptual mechanism is adaptable enough to accommodate the kinds of inter-character communication that the wolves need in order to form relationships with each other.

### 3.2.6 Emotion

In order to simulate emotional virtual characters, it is 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.

Various researchers have implemented versions of the dimensional approach; for example, Breazeal's emotional system in Kismet [Breazeal 2000] made affective assessments in a 3-dimensional space (Arousal, Valence, Stance). These assessments in turn trigger emotion elicitors, which lead to active emotion responses of joy, anger, disgust, fear, sorrow, surprise, interest, boredom or calm. The net affect that results in turn influences the robot's facial expression.

Yoon [Yoon 2000a] also featured a dimensional emotional model. She created affective synthetic characters using Russell's three axis model [Russell 1980] of Stance, Valence and Arousal. She described how her affect system could represent a six element emotional system as well by means of a network with two layers and a mapping between the three basic nodes in the top layer and the six nodes in the second layer. While Yoon used an affective tag mechanism for building synthetic characters within our group's framework, Yoon's characters did not form long-term relationships that were different with different individuals. In addition, I chose a slightly different emotional model, which replaces Stance with Dominance, since the term "Dominance" is referred to throughout the literature of natural wolf studies.

Velasquez [Velasquez 1998] based his implementation on a model drawn from several theorists [Izard 1991; Ekman 1992; Johnson-Laird 1992]. Two projects arose from his Cathexis model – Simón the Toddler and Yuppy the Emotional Robot Pet. Other researchers (e.g., [Gadanhó 1998]) have also implemented categorical models.

Various other researchers have explored the way in which computational models of emotion may affect synthetic behavior (e.g., [Cañamero 1997]). 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].

### 3.2.7 Learning

Models of computer learning abound. While the model of social relationship formation presented in this thesis is meant to integrate with the learning system currently in place in the Synthetic Characters Toolkit, the length and breadth of machine learning is outside the scope of this thesis. For an overview of machine learning, the reader is directed to [Ballard 1997] and [Sutton 1998]. The model in this document will integrate with the system described in [Blumberg 2001 (to appear)].

Nevertheless, the virtual wolves described in this thesis remember emotional relationships that they form during their young lives; as such, their social relationships exhibit a simple kind of learning. Other researchers have implemented models of emotional learning or memory. For example Yoon's "affective tags" [Yoon 2000b] allowed her characters to remember associations between locations and positive or negative experiences. Velasquez's "emotional memories" [Velasquez 1998] mark specific stimuli that resulted in strong emotional reactions so that the next time the agent encounters one of those stimuli, it re-experiences that emotional state.

Gadanhó and Hallam used an emotional system to enhance reinforcement learning in a simulated robot [Gadanhó 1998]. Balkenius and Moren [Balkenius 2000] discuss how a model of context enhances their computational model of classical conditioning and habituation. Kitano [Kitano 1995] presented a hormone-like model for affecting the learning process.

The model presented here is the first time I am aware of that emotional memories have been applied to social relationship formation.

### 3.2.8 Development

The process of behavioral and physiological development is beginning to be seen as an important part of the process of making intelligent computer systems [Weng 2001]. Various researchers have used models of development for a range of purposes. For example, researchers have used a model of development to influence the mapping from genotypic space to phenotypic space in their mechanism for evolving autonomous agents [Dellaert 1994; Hart 1995]. Developmental models have been used to build neural networks [Vaario 1991; Gruau 1993]. Di Paolo explored the role of parental influence on phenotypic development [Di Paolo 1999]. The video game *Creatures* (see Figure 3-15) featured a developmental model [Grand 1997]. Cañamero's Abbotts project also presents a developmental perspective on virtual entities [Cañamero 1997]. This listing offers a few representative examples of the ways in which developmental models are used in computational systems.



**Figure 3-15: A Norn from the game *Creatures*.**

Breazeal's Kismet [Breazeal 2000] also bears mentioning in this Development section, since that project featured a system that might provide the scaffolding for the control system of a robot or virtual entity to learn to interact with its world over a developmental process.

### 3.2.9 Social Human-Computer Interaction

Since a significant portion of the evaluation of this thesis rests on the wolves' success in interacting with human participants, it is relevant to consider the field of social human-computer interaction. The popular appeal of the AlphaWolf installation relies on people's willingness to interact with the virtual wolves as social entities.

Since 1986, Byron Reeves and Clifford Nass have led a research effort at Stanford University that has done a wide range of experiments to explore social interactions between people and

technology (e.g., [Nass 1994; Nass 1996; Reeves 1996]). Their work suggests that people will interact with technology in ways that resemble the ways in which they interact with other people.

The basic structure of the experiments done by Reeves, Nass and their colleagues is as follows: they find an experiment that has been done by social scientists about how people interact with each other; they replace one of the people in the experiment with a computer program or other bit of technology; and they re-run the experiment to see if they achieve the same results that the original experimenters found. Their experiments sweep across a wide variety of phenomena: politeness, performance-evaluation, personality, similarity-to-self, memory, team-building and gender roles, to pick out a few. Their research results are summed up by their “media equation”, that “people’s responses to media are fundamentally social and natural.” [Reeves 1996, p. 251]

Our system for controlling semi-autonomous characters is drawn from a variety of sources – virtual characters and agents, computer games, computational models of emotion, and the behavior of wild gray wolves. Our research is inspired by many of the projects below, and incorporates elements of several; nevertheless, no project that we are aware of has made a distinction between user-controlled action and autonomous emotion.

### 3.2.10 Directable Characters and Agents

Various researchers have developed mechanisms for the high-level direction of virtual characters. Hayes-Roth *et al.* [Hayes-Roth 1995] have explored “directed improvisation” as a way for users to direct and constrain the behavior of computer characters. Johnson *et al.* [Johnson 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 [Blumberg 1995] integrated autonomy with directability using a multi-level approach. The Improv system of Perlin and Goldberg [Perlin 1996] addresses the creation of believable synthetic actors, using procedural techniques to create layered, non-repetitive motions and transitions. Johnson [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; Cassell 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 1997; Magnenat-Thalmann 1998]). Bates and his colleagues (e.g., [Bates 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.

Various researchers have explored ways of expressively controlling the bodies and faces of animated characters (e.g., [Hodgins 1997; Brand 1999; Chi 2000]). Rose’s research [Rose 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 problems similar to those that we confronted in making AlphaWolf. The problem of making an autonomous creature which 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., [Chen 2001; Scerri 2001]). No research project that we have encountered uses our action/emotion distinction for separating user-control from autonomy.

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 in this paper is similar to the goals of the projects described above, our approach to this problem is novel.

### 3.2.11 Computer Games

One of the challenges of computer games is to create the illusion of life in their characters [Thomas 1981]. A relevant way in which characters may manifest life-like qualities is by having expressive social interactions with each other. *The Sims* is perhaps the most prominent of the games that feature social interactions. (See Figure 3-16.)



Figure 3-16: *The Sims*.

Many role playing games (e.g., *Everquest*, *Baldur's Gate*) allow players to take on the role of a character and interact with other non-player characters in the game (both those controlled by other users and those controlled by autonomous behavior systems). However, the interactions in these games are either scripted (in the autonomous case), or completely controlled by the human users. In *AlphaWolf*, the semi-autonomous wolves have their actions controlled by the participant, but their emotional states are controlled by the social relationship mechanism of this thesis.

The challenge of making directable characters with strong personalities is relevant in a very practical way in the making of computer games. Many games allow the user to play the role of a character exploring a virtual world – for example, Lara Croft in the game *Tomb Raider*. In many of these games, things that happen affect how the player's character is able to interact with its world; in *Resident Evil Code Veronica X*, for example, the player's character limps and moves more slowly when injured. Other games, for example *Black & White*, also offer characters who are largely autonomous but can be directed by the player.

A game called *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 our mechanism, but explores a similar action/emotion split.

### 3.3 *Summary*

To summarize the work that is most immediately relevant to this research project: the mechanism described in this thesis combines Damasio's Somatic Marker Hypothesis [Damasio 1994] with Mehrabian and Russell's dimensional model of emotions [Mehrabian 1974] to create social relationships in a pack of virtual wolves. These relationships are mechanisms by which context is preserved among different dyadic pairs of wolves [Cohen 1999]. The behavior of the virtual wolves is informed by the study of real wolves (e.g., [Fox 1971; Klinghammer 1985; Mech 1998]), and by previous work in synthetic social systems (e.g., [Hemelrijk 1996]). The social relationship mechanism is situated in the Synthetic Characters Group's Toolkit, described in [Blumberg 2001 (to appear)], [Burke 2001], and [Isla 2001b]. In the evaluation section, the mechanism is compared to other mechanisms by which entities may interact, in particular those described in [Vaughan 2000]. In addition, the AlphaWolf installation is evaluated with similar criteria to those used by other believable social and emotional agent systems – e.g., [Perlin 1996; Reeves 1996; Reilly 1996; Cassell 1999; Breazeal 2000]. Inspired by Craig Reynolds' Boids [Reynolds 1987], I have tried to keep the mechanism as simple as possible.

None of the systems that I have encountered uses an emotional memory mechanism to enable long-term social relationships among computational entities.

## 4 Implementation

In order to demonstrate that the ideas described in this thesis are viable, we created several implementations incorporating them. AlphaWolf is the most prominent of the implementations, and is described in greater depth in an earlier chapter of this thesis. In addition, I adapted the AlphaWolf code base to serve as the platform for the user tests and simulations described in the Evaluation section below. To summarize these other implementations: the user tests involved subjects interacting with the AlphaWolf installation under controlled conditions and directing the gray pup to form prescribed relationships with the white pup and the black pup, and the simulations involved a pack of six wolves feeding on a virtual carcass.

There were certain differences among the three implementations; AlphaWolf, for example, placed a premium on a full interactive experience for several networked participants, while the user tests and synthetic experiments needed to be more specific in order to target certain results. This section described the system that lies behind the implementations, and points out differences among the three specific cases.

The social relationship system is built to be part of the Synthetic Characters code base, and incorporates a variety of existing systems and parts of systems from that code base. The work described here entailed modifying several of these systems and enabling them to work together to create social relationships. In addition, the core of the research – the Context Specific Emotional Memory (CSEM) mechanism – was created “from scratch” for the social character-based installations.

The bulk of the Synthetic Characters code base is written in Java. All parts of the implementations described here are written in Java as well.

In the course of developing the three implementations, we relied heavily on various visualizers to help us understand and control the various elements of the characters and their relationships. Prior installations by our research group featured some of the visualizers we used to interpret the wolves’ social behavior. In addition, we created a new visualizer for the CSEMs to help us see how the relationships among the wolves were changing over time.

Over the next several sections, I describe the elements of our system that are integral to the AlphaWolf social relationship mechanism, and point out which parts are new to this implementation. In addition, most sections include descriptions of the differences among the three implementations and of the visualizer that helped us comprehend that bit of code.

### 4.1 Hardware

The AlphaWolf installation runs on five 933MHz single processor PIII PC’s and one dual 600MHz PIII PC, each with a GeForce 3 graphics card and running Windows 2000. In addition, there was an assortment of supporting hardware, described below. Figure 4-1 provides a diagram of the various components.

A 933MHz machine named Vole is the central behavior machine. This machine performs all of the behavior and motor calculations for all of the wolves. It displays a relatively wide camera view of the wolf pack on a 52-inch plasma display.

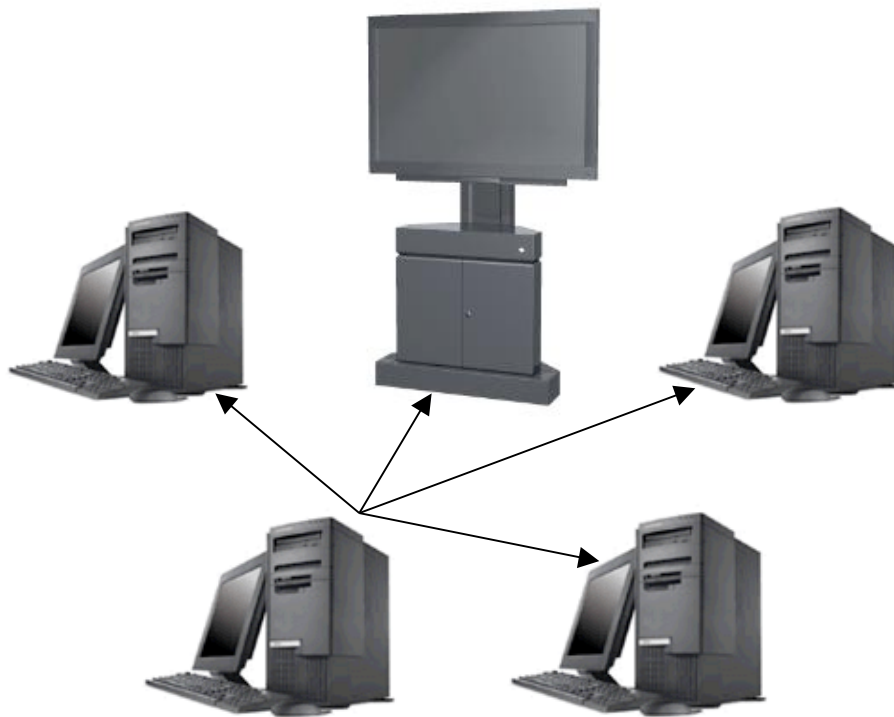
A 933MHz machine named Grebe is the audio input processing machine. It has a high-end RME 24-bit sound card so that it can accept and process inputs from three Shure SM94 microphones. It



connects to Vole via a local area network to communicate the results of its Howl Interface processing.

Three additional 933MHz machines (Bass, Cougar and Platypus) each run an abridged version of the main installation, with a full graphics system that runs a copy of Vole's motor instructions that has been sent over the network. Each of these three machines has a camera system that focuses on one pup – gray, black, or white.

The dual 600MHz machine named Silversides is responsible for producing the sound effects. The main behavior program running on Vole tells a program on Platypus what sound effects to generate and where they should be located. Platypus, in turn, sends MIDI notes to a Korg Triton synthesizer, which plays the appropriate samples to the correct channels at the right volume and pitch. The synthesizer connects to two amplifiers, each with a pair of speakers, such that the wolves may appear out of any of the four speakers.



**Figure 4-1: The hardware behind the full AlphaWolf installation.**

## ***4.2 The World***

The virtual world inhabited by the AlphaWolves is quite simple. The wolves live on a horizontal plane, defined only by the shadows they cast on it.

The trees that appear on screen are invisible to the wolves, and serve only to provide a sense of space for users. As well as setting the scene in a snowy, sparse forest, they help orient users by giving a standard reference frame when the camera moves.

In addition to the wolves themselves, there are a few invisible dummy objects, which move randomly, and to which a bored wolf might occasionally attend. These dummy objects represent leaves, sticks, or interesting smells that catch the wolves' fancy. Their presence serve the purpose

of keeping the wolves somewhat spread out. Without them, wolves have nothing to attend to but each other and rapidly clump together.

The wolves' world maintains a simple representation of time. Approximately 20 times each real-world second, each wolf has a chance to determine its behavior. While real wolves live in a more parallelized universe, the relatively high frame rate of the system keeps the virtual wolves from suffering any order-of-execution effects.



**Figure 4-2: The gray wolf walks around the virtual forest.**

The world and the creatures in it are displayed using a “charcoal renderer” written by Marc Downie (see Figure 4-2). The renderer uses a technique based on programmable vertex shaders, causing a texture map to be applied to each vertex based on the angle between its normal and the direction of the camera.

Wolves cast shadows as well, which are rendered by duplicating the skin, rendering it with a uniform dark gray color, adding the height (Z axis) of each vertex to its position on the X axis, and setting the height to be 0. This process causes the sun to appear to be situated at a 45 degree angle to the wolves, casting strong shadows that greatly enhance the three-dimensionality of the wolves.


#### **4.2.1 The Three Implementations**

Each of the three implementations occurs in a similar world. The frame rate varies a bit, based on the number of wolves in the installation (AlphaWolf had six wolves, for example, while the user tests had three), but otherwise the worlds are nearly identical. The synthetic experiments features one additional element – a single “carcass” object, on which the wolves take turns feeding.

### **4.3 The Wolves**

The virtual wolves exist in simulated three-dimensional space, moving around on the ground plane. Each wolf has 53 rotational degrees of freedom that make up its spine, legs, tail, head and face (see Figure 4-3). These nodes control the motion of each vertex of the wolf’s skin (~4000 polygons), with up to three nodes contributing to the position of each vertex.

Each wolf is composed of a variety of systems that give it certain kinds of functionality. The Action System causes the wolf to decide what to attend to and how to behave (see section 4.4). The Navigation System causes the wolf to approach its object of attention, if appropriate, to orient towards it, and to avoid collisions that sometimes occur (see section 4.5). The Autonomic Variable System maintains its internal emotions and motivations (see section 4.6). The Sensory System, Proprioception System and Perception System allow it to take in information from its surroundings and internal state (see section 4.7). The Working Memory System allows the wolf to keep track of other entities in the world over time. The CSEM System connects the Working Memory System with the Autonomic Variable System, thereby maintaining the emotional memories that form the basis of the wolves' social relationships (see section 4.8). The Morph System controls the physical growth of the wolves from pups to adult (see section 4.9). The Motor System converts the actions, emotions, and developmental elements of the wolves into concrete, real-time rotational information for the 53 nodes in each wolf (see section 4.10).



**Figure 4-3: The pup's skeleton and skin.**

These systems represent the collective work of the Synthetic Characters Group over the last five years. While each of these systems performs a wide variety of complex tasks, this chapter focuses mainly on those elements that are directly related to the social relationships of the wolves.

#### ***4.4 Action Selection***

The wolves featured the same essential mechanism of action selection as was used in the Synthetic Characters Group's sheep|dog installation at E3 in 2001 [Burke 2001], and which was mentioned in the Related Work section of this document. This mechanism involves codifying rules of behavior into ActionTuples (see Figure 4-4). Each ActionTuple features a TriggerContext, which determines the set of conditions in which the ActionTuple becomes eligible for selection, a DoUntilContext, which determines when the ActionTuple should stop executing, the Action itself, which is a call to the motor system or to another set of actions, and an (optional) object, which specifies the thing towards which the wolf should perform this action.

TriggerContext	Object	Action	DoUntilContext
Value			

**Figure 4-4: An ActionTuple.**

Each ActionTuple is a member of an ActionGroup, from which only one ActionTuple will be active at the same time. There are two kinds of ActionTuples – regular ActionTuples and startle ActionTuples. Under normal conditions, each ActionGroup chooses probabilistically from among those regular ActionTuples whose triggers are non-zero, based on the value of each ActionTuple. Once an ActionTuple has been chosen, it will stay active until its DoUntilContext is satisfied, or until the world changes significantly (defined as some non-active ActionTuple evaluating to at least double the trigger value that it had when the currently active ActionTuple became active, and at least half the value of the currently active ActionTuple). Startle ActionTuples are used for extraordinary situations (for example, something that might surprise the wolf, such as being bitten); if there is a non-zero valued startle ActionTuple, the most highly valued startle ActionTuple wins, regardless of the regular ActionTuples' values.

User interaction was modeled through startle ActionTuples, since we decided the wolves should react promptly to user input. The system of user inputs is discussed more fully below (see section 4.12). With regard to the action selection system, the user could provide inputs to startle ActionTuples that “forced” the pup to perform certain actions or to attend to certain other wolves. If a pup had not received any user interaction for some time (~5 seconds if interacting with a social partner or ~15 seconds if not interacting), it would return to the autonomous regime described below.

In the wolves' behavior system, we had a clear separation between the attention selection mechanism and the action selection mechanism. At each tick, a wolf figures out what it should be attending to, and then decides what it should do to that object of attention. This separation skirts the optional object component of the ActionTuple, but helped us craft a system in which we could address attention and action separately.

Elements that influenced the selection of action or attention were: age, emotional or motivational state, presence of some perceived context (e.g., my object of attention is within a certain distance of me), other kinds of internal state (e.g., I just switched my object of attention), and user input (e.g., my user howled into his microphone).

An example ActionTuple might look something like this:

```
//create a new ActionTuple with value 10
tuple = new ActionTuple(10);

//only do this for adults
tuple.addTrigger(amAnAdultContext);

//trigger when you're awake
tuple.addTrigger(notAsleepContext);

//trigger if you're feeling dominant
tuple.addTrigger(dominanceHighContext);

//trigger if your OOA isn't submitting
tuple.addTrigger(objectOfAttentionNotSubmittingContext);

//only activate the ActionTuple if all triggers are high
tuple.setTriggerPolicy(MULTIPLY);

//stop when your OOA submits
tuple.addDoUntil(objectOfAttentionSubmittingContext);
```

```

//stop if you feel submissive
tuple.addDoUntil(dominanceLowContext);

//do this motor skill when active
tuple.addAction(MotorAction("DOMINATE"));

//Do it to your OOA
tuple.addObject(objectOfAttentionContext);

```

This ActionTuple causes an awake, dominant, adult wolf to perform a DOMINATE motor action when confronted by an individual who is doing anything but submitting. The wolf will persist in dominating until the object of attention submits, until the wolf itself begins to feel submissive, or until a startle ActionTuple interrupts it. ActionTuples like this one, coupled with a perception system and a dynamic emotional state, form the basis for social action selection of the wolves. In the Emotion and CSEM sections below, I discuss how the emotional state is determined. In the Perception section, I discuss how a wolf perceives what its object of attention is doing.

The set of dominance behaviors is more complex than this example might lead one to believe. While the core structure of the DOMINATE behavior reflects the description above, the DOMINATE motor action is a compound set of skills – for example “stand with tail and ears erect”, “bite at”, or “growl.” These individual components are combined by a variety of means, thereby serving as a simple, two-level behavior hierarchy. While the wolves’ behavioral repertoire is simple enough to be modeled by this two-level hierarchy, more complex creatures will demand a full hierarchical action-selection mechanism.

Just as there is an ActionTuple and second-level set of sub-behaviors for DOMINATE, there are corresponding action-selection structures for ACTIVE\_SUBMIT and PASSIVE\_SUBMIT. Both of these behaviors have low dominance as a trigger. ACTIVE\_SUBMIT occurs when a submissive wolf initiates an interaction with a wolf, and that wolf is not attempting to dominate the wolf. PASSIVE\_SUBMISSION occurs when the submissive wolf is the target of its social partner’s dominance display. PASSIVE\_SUBMISSION therefore occurs in two main contexts – when a dominant wolf approaches it and “demands” submission, or when it is ACTIVE\_SUBMITTING to a wolf, and that wolf decides to attend to it and DOMINATE it. Both of these contexts are represented identically – “myObjectOfAttentionIsDominatingMe” – hence, there need be only one ACTIVE\_SUBMIT ActionTuple.

In addition to the core social dominance behaviors, the wolves have an assortment of other behaviors that flesh out their repertoire.

For example, wolves will howl if their users direct them to do so, or under certain other conditions. Howling serves the purpose of making the other members of the pack aware of an individual’s position. If a wolf wants to find that individual, it can locate it by means of the sound, even if the wolf is not visible through the fog in the world.

Wolves sometimes wander around if there are no immediately pressing social interactions, attending instead to one of the dummy objects mentioned in the World section above. This mechanism for dispersing the wolves has proven to be an important element of visually clear social relationships, since it is harder for a participant to tell who is dominating whom if three or more wolves are all in the same place and interacting with each other.

Wolves sleep if they are fatigued. Fatigue gradually increases until a pup is compelled to sleep. User interaction reduces fatigue, so a pup will never sleep if it is being actively directed.

If two wolves come together and neither of them has a clear dominance agenda, they may instead play with each other. Play behavior appears to serve quite a variety of purposes in wolf development [Fox 1971; Bekoff 1974], and in social development in particular. In the AlphaWolves, play behavior is modeled as a separate set of ActionTuples from the dominance and submission behaviors, but is nevertheless tied into the overall behavioral scheme since it causes wolves to pair off, which often leads to social dominance interactions.

Finally, the wolves have a simple set of reflexes that help them cope with one of the hard problems in simulating social behavior – close contact. If, for example, one wolf bites at another wolf, the bitten wolf should flinch away from the bite. In real wolves, this is the result of a complex motor-learning problem that is beyond the scope of this document and the AlphaWolf project. We implemented reflexes as a set of very high priority behaviors that match certain behaviors taken by the social partner. In the biting example above, each of several bites has a matched “flinch” behavior that causes the target wolf to perform a realistic avoidance of that bite. The successful coordination of these behaviors is predicated on the success of the navigation system, discussed in the next section, which causes two interacting wolves to be loosely facing each other.

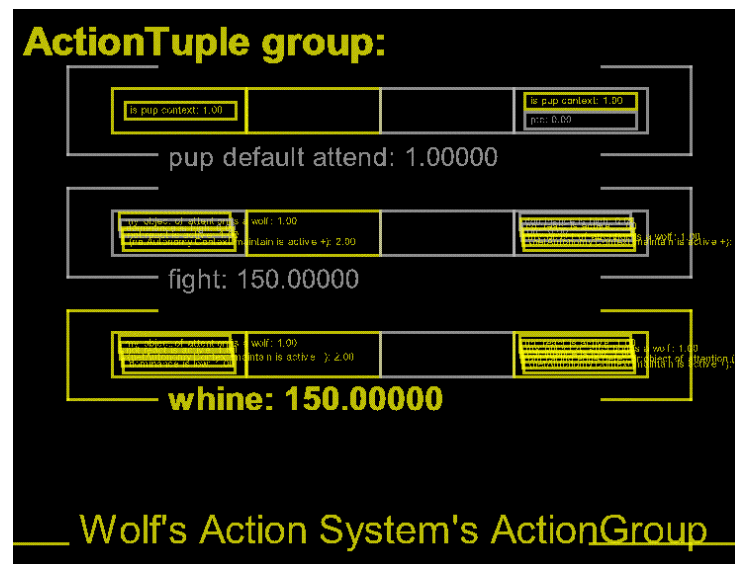


Figure 4-5: The Action System visualizer for a wolf in the act of whining.

#### 4.4.1 Visualizer

The Action Group visualizer that we use in AlphaWolf’s action selection was developed by Robert Burke, and reimplemented by Matt Berlin. In the visualizer, each ActionTuple is represented by a horizontally arranged set of joined rectangles, representing TriggerContext, Object, Action, and DoUntilContexts. Within the Trigger, Object and DoUntil rectangles, all of the various relevant contexts are stacked. All inactive elements are gray, while all active elements are yellow. This color difference makes it easy to pick out the active actions, or to determine why a certain element is active or inactive due to the color of its constituent parts. In addition, the mouse can be used to pan and zoom in on the various bits of the visualizer.

The attention selection mechanism and action selection mechanism are visualized in two distinct windows.

#### **4.4.2 The Three Implementations**

The action selection mechanism in the user tests started from that of the AlphaWolf installation. In order to make the directions fool-proof, it was simplified significantly. The play behaviors were removed, as was howling. Development was removed, as were the adults, so that the subjects experienced three young pups who never age.

While the AlphaWolf installation was a three person experience, the user tests had only one interactor at a time. Therefore, the autonomy of the other two pups was increased (since they would have no user input to keep them interesting).

The simulated experiments started from the user test demo, but showed certain differences. The most prominent of these modifications was a small suite of behaviors that allowed the wolves to approach and eat a carcass. In addition, certain values needed to be tweaked to accommodate the presence of six boisterous pups rather than just three. For example, the initial behavior of a group of six pups was to clump together, since even a pup desiring to wander away would usually be dragged into an interaction by one of its five litter mates. Turning down the desire to interact relative to the desire to wander away caused the larger litter to interact frequently but not constantly.

The AlphaWolf action selection mechanism was primarily created by the author, Matt Berlin, Derek Lyons and Jennie Cochran. The action selection mechanism for the user tests and synthetic experiments was adapted by the author.

#### **4.5 Navigation, Orientation and Collision**

In addition to the main action- and attention-selection mechanisms, there are several systems for navigation, orientation and collision-avoidance that form the underpinnings for the wolves social relationships. Without the ability to find each other, face each other, and avoid (or at least correct) any collisions, the wolves would not be able to engage in graphically compelling social interactions. In addition, the problems that arise in confronting this level of social encounter might shed some light on issues relevant to biological social behavior that are not addressed in a less complex simulation.

The first component of this group of systems is the ability to move from one place to a specific other place. This ability lets wolves approach each other. To do this, wolves perceive the position of their target, and then WALK or GALLOP there. Both WALK and GALLOP are blended motor skills with left-hand, straight, and right-hand examples. The left and right examples define the maximum turning degree that the wolf can exhibit while locomoting. The Navigation system calculates how much of each of the three examples to blend together to get the desired walk angle. So, for example, if a wolf wanted to run to another wolf ninety degrees to its left and far in the distance, it would start by running sharply left, and then gradually blend into straight run cycles as it came to face straight toward its target.

Wolves choose which gait to employ based on several factors. Wolves generally gallop, except in the following cases: 1) the action that their user has requested them to take does not match their current emotional state (e.g., a submissive wolf is directed to growl) or 2) the wolf is sufficiently close to another wolf that galloping would make it likely for them to interpenetrate.

The wolf may transition directly from WALK to GALLOP and back. This ability was especially useful in the final stages of approaching an object, because it allowed the wolf to gallop to within a certain distance and then walk the last few feet. By preventing a wolf from galloping within a certain distance of another wolf but allowing it to walk a bit closer, a natural-looking and relatively accurate approach behavior was achieved.

When a wolf is too close to another wolf to get correctly oriented during the navigation stage, it may need to perform some specific orienting behaviors instead. These behaviors involve turning left and right while standing in place, or while transitioning between sitting and standing. In the action selection mechanism, an ActionTuple can be set to need to be correctly oriented before executing. Dominating behaviors, for example, need to be oriented correctly; biting the wrong way makes a character look very stupid. To orient, the wolf will turn left or right as much as it needs to before beginning the action it has chosen to do.

Because social contact is strongly correlated with close proximity, we have tried to make it possible for the wolves to get close to each other. While this nearness makes possible some of the most powerful moments between the wolves, in the absence of real-world physics it also sometimes leads to collisions between wolves. Nothing breaks a viewer's suspension of disbelief like one wolf crossing *through* another wolf, rather than going around it. And if avoiding static objects is a challenging but tenable task for an animated creature, having to avoid other animated creatures running in real time is even more problematic. This problem will be addressed further in the Future Work chapter.

We implemented a variety of ways for wolves to avoid interpenetrations and to correct them if they happen. Essentially, wolves have the ability to step forward, backward or to either side. When they detect that they are in the "personal space" vector field of another wolf, they choose whichever action will get them out of that vector field as rapidly as possible. This vector field is weighted slightly toward the heads of the wolves, to take into account the fact that wolves move forward most often under normal circumstances.

The navigation system is implemented using the same ActionTuple framework as the action-selection mechanism. The various Actions of the navigation system compete probabilistically based on their TriggerContexts, and run until their DoUntils are satisfied. The navigation system runs *after* the action system, and therefore has priority to change the behavior that is ultimately executed by the motor system. For example, if a wolf wants to dominate another wolf who is some distance away, the navigation system will overwrite the DOMINATE motor action with a GALLOP motor action until the wolf is close enough to have a convincing interaction.

### 4.5.1 Visualizer

Since the navigation system, like the action system, uses ActionTuples, it also employs the same style of visualizers. The navigation ActionTuples appear in a separate window from the action system ActionTuples, denoting that they are members of a separated ActionGroup.

### 4.5.2 The Three Implementations

There were only a few differences in this section among the three projects. First, advances in the orientation mechanism, which allow a creature to orient both before and after it navigates, came online after AlphaWolf but were included in the user tests and synthetic experiments. In addition, the approach distance (how close a wolf gets to its object of attention before it stops) was increased in the user tests, since I wanted subjects to be as undistracted by collision as possible,



even if the increased inter-character distance slightly compromised the believability of the social interactions.

## 4.6 *Emotion and Motivation*

The first main component of the AlphaWolf social relationship mechanism itself is the emotion system. The relationship mechanism involves learning associations between the individual's current emotional state and the presence of another individual; therefore, a dynamic emotional state is central to the relationship-building process. While this section is the only one with the word "Emotion" in its title, it is not the sole section to deal with some element of the topic; emotion features prominently in several other sections of this chapter as well (in particular, the sections on CSEMs and Expression).

The model of emotion in the AlphaWolf installation as presented at SIGGRAPH is very simple – a single floating-point value called *dominance*, which varies from 0.0 to 1.0. In the Synthetic Characters code base, this is represented as an `AutonomicVariable`, which has a variety of convenience functions for specifying drifts, set points, etc.

Each wolf's dominance value is affected directly by its interactions with the world and indirectly by its memories. 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. Also, encountering an individual whom the wolf remembers is dominant to it will make the wolf's own dominance decrease, and meeting a historically submissive wolf makes the wolf's dominance increase.

In addition, in AlphaWolf, the actions of the wolf itself changed its own emotional state. So, for example, a growling wolf's dominance would gradually climb, albeit at a slower rate than if another wolf were submitting to it. In the user tests and synthetic experiments, I removed this functionality for simplicity's sake, and the relationships remained clear.

The update algorithm for the emotional state of the wolf is as follows:

$$E' = E + \text{drift} + \sum \text{sideEffects}$$

**Equation 4-1: The emotional update algorithm.**

In this equation,  $E'$  is the new emotional state,  $E$  is the emotional state on the previous tick, drift is the amount the emotion will move toward its set point each tick, and sideEffects are the result of phenomena impinging upon this emotion (for example, being growled at is a sideEffect that reduces dominance).

The emotional state has an impact on the wolves' behavior in three main ways – through emotional memories, action selection, and expressiveness. The next three paragraphs describe each of these in more detail.

Most importantly to this thesis, the emotional state serves as the basis for the emotional memories that the social wolves form of each other. An individual's dominance at the moment when an interaction with a social partner comes to an end affects the emotional memory that the individual maintains about the partner. The exact mechanism by which this memory revision happens is described in the CSEM section below.

Second, when a wolf is behaving autonomously (i.e. not under a user's control), its emotional state affects what actions it chooses to take towards its social partners. For example, a submissive

wolf might choose to roll over on its back, while a dominant wolf might growl. Emotions directly feed into the action selection mechanism by means of `AutonomicVariableContexts`, which act as `TriggerContexts` to the relevant `ActionTuples`.

Third, the wolf's emotional state affects the style in which the wolf takes its actions. For example, a wolf with a current dominance value near 0.0 might walk with its tail between its legs and its ears back, while a wolf with a higher dominance value will hold its tail and ears erect as it walks. The expressive component of the wolves' emotional states is discussed in greater depth in the Expression section below.

#### 4.6.1 Visualizer

The visualizer that we use for emotions such as dominance is a simple slider bar that shows the current state of the variable (see Figure 4-6). An important piece of functionality in this slider is the ability for a person to set the emotion's value in real time as well as just looking at its value. By dragging the slider, the programmer can directly affect the wolf's emotional state. This functionality has proved to be invaluable in reliably creating specific scenarios for developing and debugging.




Figure 4-6: These visualizers show that Wolf5 has a very low dominance and a mid-range hunger.

#### 4.6.2 The Three Implementations

In addition to the key emotion of dominance, which features prominently in all three implementations, the wolves in the AlphaWolf installation have two additional motivational states that affect their behavior – boredom and fatigue. A high level of boredom causes a wolf who has not received any input from its user in some time to begin to behave autonomously. Boredom drifts up continually, and drops to zero when the user interacts in any way. Fatigue serves to cause the pups to start off asleep, and periodically to go back to sleep if the users do not interject some activity into the scene. Real wolves, after all, spend a large portion of their days asleep.

The simulations also had an additional simple interface by which a person or an automatic system could force a pup to be dominant for a specified duration.

### 4.7 Sensation and Perception

Synthetic perception is central to all of the characters that we have built in the Synthetic Characters Group, and is particularly important to their social competence. In order to be able to form relationships with specific individuals, a wolf needs to be able to detect the presence of those individuals, to tell them apart from one another, and attend to them. Wolves also need to be able to detect the appropriate portion of the wolf to which to attend (e.g., to be able to actively submit at the dominant wolf's face, rather than any other body part). Finally, wolves need to be able to discern certain information about their social partners (e.g., is this wolf attending to me).

These kinds of skills are implemented through the Sensory System and the Perception System. In this section, I provide a cursory overview of these systems. For a full treatment of these systems, please see [Burke 2001; Isla 2001a].

There are two main parts of the perceptual mechanism in the wolves. First, each wolf posts certain information to a central blackboard [Isla 2002]. This information includes a unique ID, its position, its pose, its object of attention, and a dominance value. These data are bundled into a *DataRecord*, which can then be analyzed by the Sensory and Perception Systems of other wolves. While it may seem simplistic for entities to identify each other by means of a unique ID, real wolves appear to be able to distinguish each other by scent [Beaver 1999].

When a wolf perceives its environment, several things happen. First, it senses its environment, receiving *DataRecords* from the world. These data are filtered by the Sensory System according to a variety of criteria, for example, visibility angle, visibility radius and smell radius. This culling stage limits the amount of information that must be addressed more closely, since distant entities will be dismissed early in the process. (This stage wasn't very relevant in these implementations, since there were only three or six wolves in the world. Nevertheless, they would be crucial in a larger system with, for example, multiple packs of wolves on thousands of virtual square miles of tundra.)

Second, these data are evaluated by a construct called the Percept Tree. During this stage, the *DataRecords* are evaluated by individual percepts (e.g., "Is this a wolf?", "Is this making a sound?") The Percept Tree is a hierarchically organized structure that includes everything that a wolf can perceive – shapes, sounds, smells, conjunctions (e.g., "This is my object of attention and it's submitting.") The Percepts form the basis of many of the *TriggerContexts*, *DoUntils*, and objects of the *ActionTuples* in the wolves' behavior systems.

An interesting aside about the Percept Tree is that for most of the development of *AlphaWolf* it was nearly flat. During this time it was one of the most computationally expensive parts of our system, since every entity was evaluated by every Percept on every clock tick. Near the end of the project's development, Jesse Gray converted it into a much deeper hierarchy without sacrificing any of its functionality. Because of the culling that goes on in the hierarchy (i.e., if a Percept evaluates a *DataRecord* to 0, none of the Percept's children try to evaluate that *DataRecord*), this was the source of a significant speed-up of the installation, taking it from 12 frames per second to around 20 frames per second.

One attribute of a wolf that another wolf can detect is its object of attention. For example, a wolf can tell if its object of attention is attending to it. It appears that real animals have "custom wiring" for detecting gaze [McNeill 1998], therefore we believe that it is reasonable for the *AlphaWolves* to be able to detect this information about their social partners directly. The process by which a wolf detects another wolf's object of attention occurs through its processing of that wolf's *DataRecord*, which includes a separate field with the name of its object of attention.

The percept evaluations that occur when a *DataRecord* is pushed through the Percept Tree also result in Beliefs in the wolf's memory. Each wolf forms a Belief for each other entity that it encounters; it is through these Beliefs that wolves take action on the world. For example, the Navigation System causes the wolf to run to the remembered position of the wolf's object of attention, extracted from its Belief about that target entity.

The Perception System is also the pathway through which wolves "hurt" each other. the value posted to the blackboard by each wolf is used to modify its effect on other wolves. So, for

example, a growling wolf with a dominance of 1.0 will cause its object of attention to become more submissive each tick than one with a dominance of 0.7.

Here is an example of how this process works, shown graphically in Figure 4-7. Wolf **A** decides to growl at wolf **B**. **A** posts a DataRecord of itself (DR1) to the world's blackboard (**BB**) that might read, in English, as “**A**, at location (10, 0, 0), is growling at **B**, with dominance = 1.0.” **B**, in turn, senses this DataRecord (either on the same clock tick, if **B** runs after **A**, or on the following tick). **B**'s Sensory System (**S**) determines that **A**'s position is within range, and passes the DataRecord to the Perception System (**P**). Each Percept at the top level of the Percept Tree evaluates the DataRecord; if its evaluation is above a certain threshold, it passes the DataRecord to its children (Figure 4-10 shows the Percept Tree). Through this process, the wolf perceives information such as “something is growling in the world”, “there is a wolf near me”, or “there is something with me as its object of attention.” Because all of these Percepts were sensed at approximately the same position where the wolf remembers perceiving a certain other wolf on the previous tick (i.e., it has a Belief (**B**) that there is a wolf at that location), the percept evaluations are attributed to that Belief. Thus, the wolf's Belief at that location is now believed to be the entity that generated the growl. Through this process, the various attributes of another wolf are reassembled into a single perceived and remembered entity, on which the wolf can then act through the object of attention mechanism (**OOA**) in the action selection mechanism (**AS**). The growling-at-me affects the emotion system (**E**), causing **B**'s dominance to decrease, and the Belief of **A** therefore becomes the target of **B**'s submissive behavior, chosen by the action portion (**Act**) of the action selection system. The action causes a new DataRecord (DR2) to be posted.

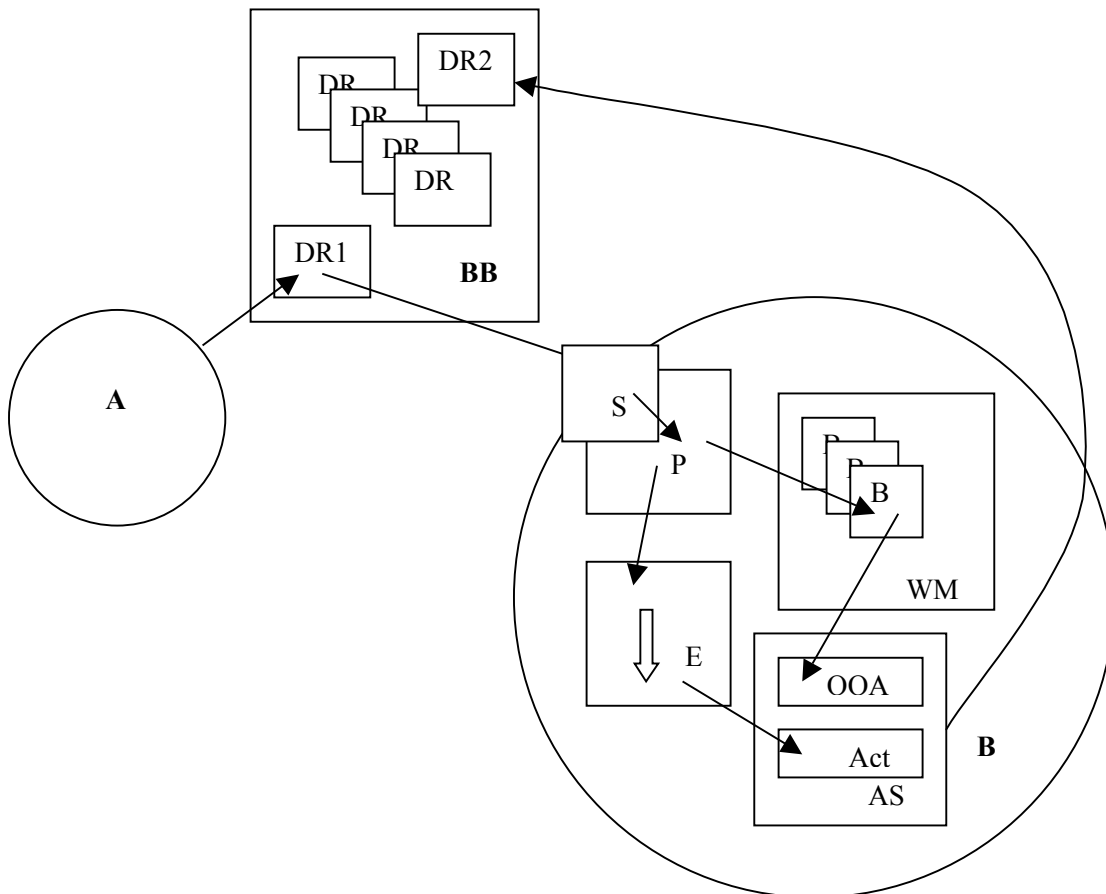


Figure 4-7: A sample interaction as perceived by the Wolf B.

### 4.7.1 To Whom To Attend

Making a system by which the wolves could effectively decide to whom to attend was surprisingly challenging. A wolf needs to be persistent in its attention, but distractible (i.e., liable to have its object of attention changed on account of an action by an individual who is not currently its object of attention) under the right set of circumstances. User-direction makes the situation even more complex.

In autonomous wolves, we settled on a system in which a wolf that has just switched its object of attention in response to user input maintains that object of attention for a certain amount of time (~4 seconds), regardless of external stimuli. (In the case of AlphaWolf, the only relevant external stimuli were other wolves. However, in a more complex world, the mechanism for maintaining a user-selected object of attention would need to be more complex.) After this duration, a wolf attending to a submitting wolf can be distracted by a dominating wolf, but not the other way around. In addition, if the wolf's boredom climbs above a certain level, it occasionally chooses a dummy object to attend to, and chooses an appropriate action (e.g., "run to and sniff", "run to and sit", etc., chosen randomly) to do to it. If another wolf interacts with it, it will switch its attention to that interactor. While enforcing the user-selected object of attention does limit a wolf's ability to react or be surprised during this short period, it prevents rapid aliasing between a user's direction and the continuing attentions of an interacting social partner. We set the duration of the lock on object of attention to be as short as possible, while still solving this problem.

In a directable wolf, any input from the user is immediately obeyed. Therefore, if the user clicks on another wolf or on the ground, the wolf immediately switches its object of attention to that item. This choice will be maintained absolutely for a short period of time, unless the user clicks again. After the enforced period has elapsed, another wolf may steal the wolf's attention by interacting with it.

The result of both of these methods of attention selection is that wolves exhibit some persistence of attention, but are nevertheless interruptible by the user, by other wolves, and occasionally by random dummy objects, in a way that feels convincing to an external observer, and that serves the goal of relationship formation appropriately.

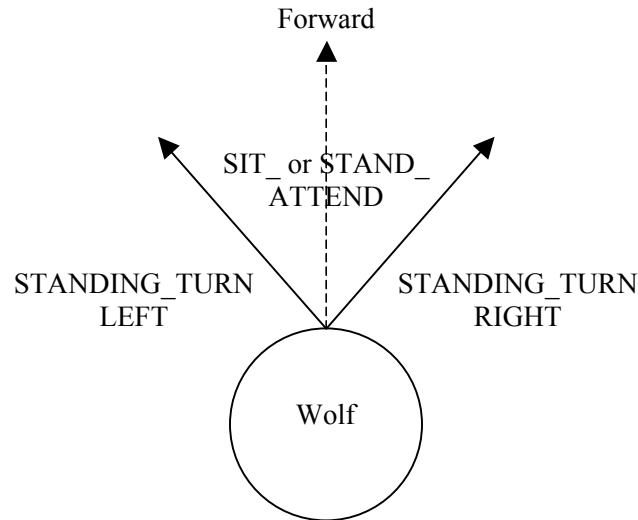
### 4.7.2 Look At

The simplest way in which a wolf acts on its object of attention is to look at it. The wolves have an intricate "look at skill", involving several blended poses and layered motor actions (see Figure 4-8). Previous characters that we have built simply looked at their objects of attention by rotating their head node; the wolves do somewhat more than this.

If a wolf is in a standing position, and is oriented such that it is facing within a certain angle toward its object of attention, it plays a STAND\_ATTEND behavior, that blends a look\_at\_right pose with a look\_at\_left pose. These poses involve some bending of the spine and neck, and some shifted weight. We then allow the head node to turn toward the object of attention, layered on top of the underlying look\_at pose. This combination has the effect of making the wolf look much more natural in a held pose looking at something. Without the full-body animation, the neck rotation often looks awkward.

If the wolf is standing but is not facing toward its object of attention, the Navigation System will cause it to orient toward its object of attention with a STANDING\_TURN animation. Once it is within a reasonable range, it will switch to STAND\_ATTEND, as described above.

If the wolf is sitting and is facing its object of attention within a certain angle, it will perform a SIT\_ATTEND, similar to the STAND\_ATTEND. If, however, its object of attention moves out of range, the wolf will perform a blended SIT\_TO\_STAND behavior during which it will orient, and will continue to perform STANDING\_TURN until it is facing its object of attention, and then sit back down into SIT\_ATTEND.



**Figure 4-8: When the wolf's object of attention is within a certain angle of its forward vector, it attends. If it is beyond this angle, it turns.**

If the wolf is performing some other action, it simply causes the head to rotate, within certain limits, to face as directly as possible toward the object of attention (see Figure 4-9). This simple layering works quite well with directional animations – for example, WALK or GALLOP. As described above, the Navigation System blends left-hand and right-hand examples to get the wolf where it needs to go; since the wolf is already turning in a certain direction, it looks natural that the head is turning that direction too.



**Figure 4-9: The gray wolf pup showing a strong look at skill.**

Although looking at things is central to the appearance of attention in the virtual wolves, the look at skill is not universally applied to the wolves' animations. Certain animations (for example, SLEEP or HOWL) would be ruined by a continuously tracking look at skill. During these animations, the look at skill is suppressed.

While our look at skill is generally quite serviceable, it does have certain problems. For example, a pup rolled over on its back in submission will sometimes have its head pass through one of its arms when it switches its object of attention. Avoiding self-penetration is among the "hard problems" of real-time motor control, because of the difficulty of detecting the penetration (bounding boxes or bounding ellipses are usually used to represent an entity during collision calculations, rather than the more computationally intensive full model), and also the difficulty of choosing a realistic-looking alternative to the motion path that would have caused the collision.

The virtual wolves, like real wolves, do not always look at their object of attention. In particular, submissive wolves often look away from the wolf dominating them, as making eye contact is a sign of dominance. The AlphaWolves do this as well, turning away from a dominant social partner and only glancing at it occasionally, as a sign of submission.

An important change that we made to the system midway through the development of AlphaWolf was to change the focus of a wolf's attention from the root node of another wolf to its head. Among non-social entities, it may not matter exactly which body part they treat as the "center" of an individual. However, as we started implementing the social behaviors of the wolves, it rapidly became apparent that attending to the pelvis was woefully inadequate for creating a compelling social interaction. It very much does matter, from the point of view of a human observer, where one social being is interacting with another. When the wolves started targeting each others' faces, the believability of their interactions became significantly better.

To summarize the role of the Sensory and Perception Systems in creating social relationships, a character is able to:

- identify individuals
- choose and maintain its object of attention in a reasonable fashion
- perceive attributes of its object of attention (e.g., position, pose, dominance)
- assess when it is the object of attention of another individual
- look at its object of attention in a believable way.

The skills described above serve two main purposes in our implementation. The first purpose is to make the wolves able to form their social relationships effectively. The second is to convey a sense of attention to the human interactors and audience members watching those relationships. By taking a biologically inspired model, these two elements work together, rather than against each other. The AlphaWolves appear to be attending to their social partners primarily because they *are* attending to them.

### 4.7.3 Visualizer

The Percept Tree has an elaborate visualizer that allows it to be examined and manipulated. The branching structure of the Tree (see Figure 4-10), combined with the large number of Percepts, results in quite a tangle of leaf nodes. The manipulation component of the visualizer allows any node to be dragged to a new location, and the surrounding nodes attempt to conform to the new configuration by repelling each other. Percepts that are active light up in yellow. Inactive

Percepts appear in gray. Unlike the emotional visualizer, manipulating the Percept Tree has no effect on the behavior or internal state of the wolf itself.

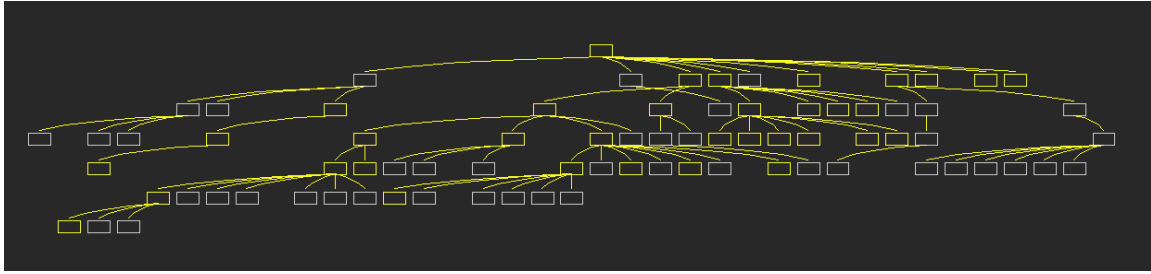


Figure 4-10: The Percept Tree.

#### 4.7.4 The Three Implementations

The structure of the Percept Trees was quite similar in each of the three implementations, although the AlphaWolf installation featured a greater number of Percepts, due to the presence of the Adults.

### 4.8 Context Specific Emotional Memories (CSEMs)

The central representation of the social relationship mechanism presented here is the Context Specific Emotional Memory (CSEM). This representation is the essence of the virtual wolves' social relationships, the remembered construct by which they keep track of their interaction histories with each other, and allow those histories to affect their current and future interactions with each other. The next few paragraphs describe exactly what happens in the CSEM mechanism with regard to two virtual wolves, **A** and **B**.

From the point of view of an individual wolf **A**, an interaction with a partner **B** begins when **A** has **B** as its object of attention, and perceives that **B** is reciprocally attending to him. For **A**, the interaction ends when it changes its object of attention so that it is no longer attending to **B**, regardless of whether **B** is still attending to **A**. In fact, **A** would have no way of telling whether **B** was still attending to it, since that information should be unavailable to it unless **B** is its object of attention.

The first time individual **A** ends an interaction with individual **B**, it forms an “emotional memory” of **B**. Our model of an emotional memory contains three pieces of information – the unique ID of **B**, an emotional value, and a confidence value.<sup>2</sup> When the emotional memory is first formed, the emotional value stores the emotion that **A** was feeling at the time its interaction with **B** ended. Since the interactions that **A** has had with **B** may have altered its emotional state over the course of their interaction, forming the emotional memory at the end of the interaction rather than at the beginning will reflect the emotional content of the relationship more accurately.

The next time **A** switches to have **B** as its object of attention, its emotional memory of **B** will influence its current emotional state in proportion to its confidence in that model. The basic formula by which the emotional memory is applied to the current emotional state is:

<sup>2</sup> Although the social relationship mechanism that we describe treats *individuals* as emotional significant stimuli, a stimulus does not have to be an individual – only a *causative entity* [Damasio 1994]. Forming emotional memories of other kinds of stimuli (e.g., the presence of two wolves at the same time) could result in other kinds of relationships (e.g., alliance formation).



$$E_{\square} = (C * E_m) + ((1 - C) * E)$$

**Equation 4-2**

where  $E_{\square}$  is the wolf's new emotional value,  $C$  is the confidence value of the emotional memory being applied (between 0.0 and 1.0),  $E$  is the wolf's emotional value prior to the application of the emotional memory, and  $E_m$  is the emotional value that is stored in the emotional memory.

The “apply” equation above instantaneously causes the full effect of the CSEM to be applied. A more complex variant of the apply method was used in the AlphaWolf installation. This variant caused the effect of Equation 4-2 to be distributed based on space and time. In space, the effect was only applied when the wolf was in close proximity to the social partner. Because of this spatial application, wolves would gradually exhibit the CSEM's effect as they ran toward social partners. In time, a limit was placed on how fast the CSEM's effect could happen – even if the wolf switched objects of attention while very close to the new object of attention, the CSEM for that wolf would be applied over 8 ticks. This gradual application prevented the wolf from “snapping” to a new emotional state.

At the end of each successive interaction, **A** revises its emotional memory of **B**. Upon revision, two of the three elements of an emotional memory are changed – confidence and emotional value. We revise confidence before we revise the emotional value so that the change in the emotional value will reflect the change in confidence, thereby preserving the effect of deviations from the expected emotional interaction.

The formula by which confidence is revised is:

$$C_{\square} = C + ((T_c - |E_m - E|) * L * (\text{Min}(C, 1 - C)))$$

**Equation 4-3**

where  $C_{\square}$  is the new confidence value for the emotional memory,  $C$  is the previous confidence,  $T_c$  is some confidence threshold between 0.0 and 1.0,  $E_m$  is the emotional value that is stored in the emotional memory,  $E$  is the wolf's current emotional state, and  $L$  is a learning rate. (Multiplying the learned component by the Min of  $C$  and  $1 - C$  effectively clamps the confidence value to between 0 and 1, and helps to polarize relationships.)  $T_c$  and  $L$  were chosen by the author to specify the ease and speed with which relationships were formed. If  $T_c$  is too low, the relationships do not readily converge; too high and the relationships become hard to change.  $L$  determines how many interactions it takes to form the relationships; in the 5 minute interaction of AlphaWolf, for example, we decided that it should take approximately three interactions with another individual to form a strong relationship with it. For a longer interaction,  $L$  should be lower so people can explore the subtlety of the behavioral repertoire.

We then revise the emotional value stored in the emotional memory:

$$E_{m\square} = (C * E_m) + ((1 - C) * E)$$

**Equation 4-4**

where  $E_{m\square}$  is the new emotional value stored in the wolf's emotional memory,  $C$  is confidence,  $E_m$  is the emotional value in the memory prior to the revision, and  $E$  is the wolf's current emotional state.

These equations represent a simple (but for our purposes perfectly serviceable) implementation of the mechanism described. For more elaborate social behavior, any or all of these equations might be made more complex.

The emotional memories described here function as remembered constructs by which an entity keeps track of its interaction history with another entity. They allow that history to affect its current and future interactions with that entity. These two elements satisfy the definition of a social relationship that we offered at the beginning of this document.

In order to watch for the beginnings and ends of interaction episodes, each wolf has an “InteractionMonitor”, which runs every tick and tests the wolf’s object of attention against its object of attention on the previous tick. If the wolf’s object of attention is the same as before, it tests to see if it is still that object of attention’s object of attention.

## 4.8.1 Visualizers

There are two main kinds of visualizers that we have used for viewing and debugging the CSEMs. The first shows all the relationships of all the pups at a specific time. The second shows the average dominance of each pup over a long time scale.

### 4.8.1.1 Instantaneous

The first visualizer, written by Jennie Cochran, is a diagram of each pup’s relationships at the current moment in time; it gets updated in real-time while the system is running. The visualizer in Figure 4-11 depicts a pack of four wolves (named “gray”, “black”, “white” and “red”), each row represents the set of relationships maintained by an individual pup. The amount of purple in each small square shows how dominant the pup feels towards the specific social partner whose name is above that square; more purple, more dominant. The first row depicts each relationship held by the red pup (hence the label “red” at the left side). The small amount of purple in each of red’s relationships shows that it feels submissive to the three other pups. The gray pup (in the last row) is the most dominant of the pups.



Figure 4-11: The visualizer of all the pups' relationships.

The amount of blue in each square shows that pup's confidence in that relationship; more blue means a less confidently held relationship. The blue is centered on the top of the purple portion (like an error bar) to demonstrate an approximation of the range of dominance that the pup expects to be feeling at the end of a future interaction with that individual.

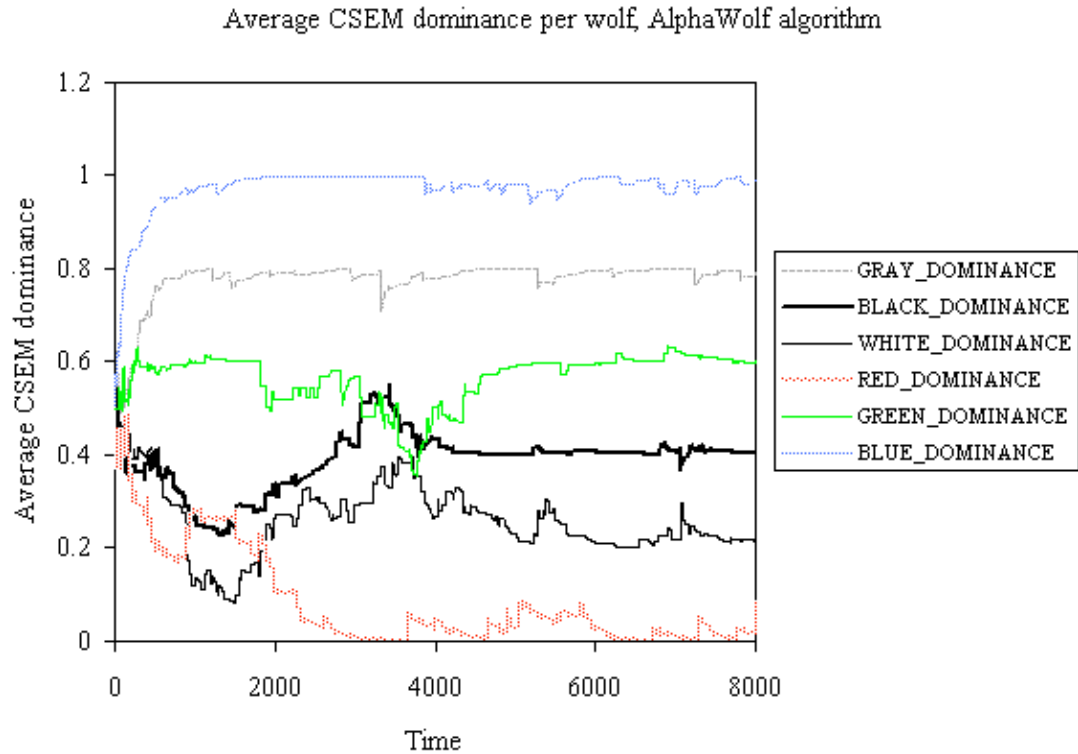
Since each of the small squares represents one pup's interaction history with a specific other pup, a dyadic relationship is essentially the sum of two of the squares. In a mature pack, most pairs of pups have well-matched relationships of each other. In Figure 4-11, for example, gray and black have a clear relationship in which gray is dominant to black. Both pups agree – black feels submissive with respect to gray, and gray feels dominant with respect to black. In fact, the only substantially contested relationship (i.e., one where neither of the two pups is confident in it) is the one between the white pup and the black pup. Their relationships with each other are in the middle of the range with regard to dominance, and both show little confidence in the accuracy of their models.

Each small square is dynamically added to the visualizer when that wolf first notices a new social partner. Therefore, the ordering of the rows, and of the specific relationships within rows, do not correlate with anything except the initial order of encounters. This attribute gives the visualizer the ability to introduce new members dynamically to the pack. We did not, however, write a mechanism by which individuals could be removed dynamically from the pack. This points to an area of future work for our system – the ability for the wolves to cull relationships that are no longer relevant (see section 7.2.1).

#### ***4.8.1.2 Continuous***

The second visualizer (see Figure 4-12) is used to show the pups' respective dominances over time. Each pup's overall dominance value is determined by taking the average of the dominance values in each of its CSEMs. This value gives a good overall view of the pup's relationships with its litter mates; for example, a wolf who feels completely dominant around all of its pack mates would have an average dominance of 1.0, while a wolf who is universally submissive would have an average dominance of 0.0.

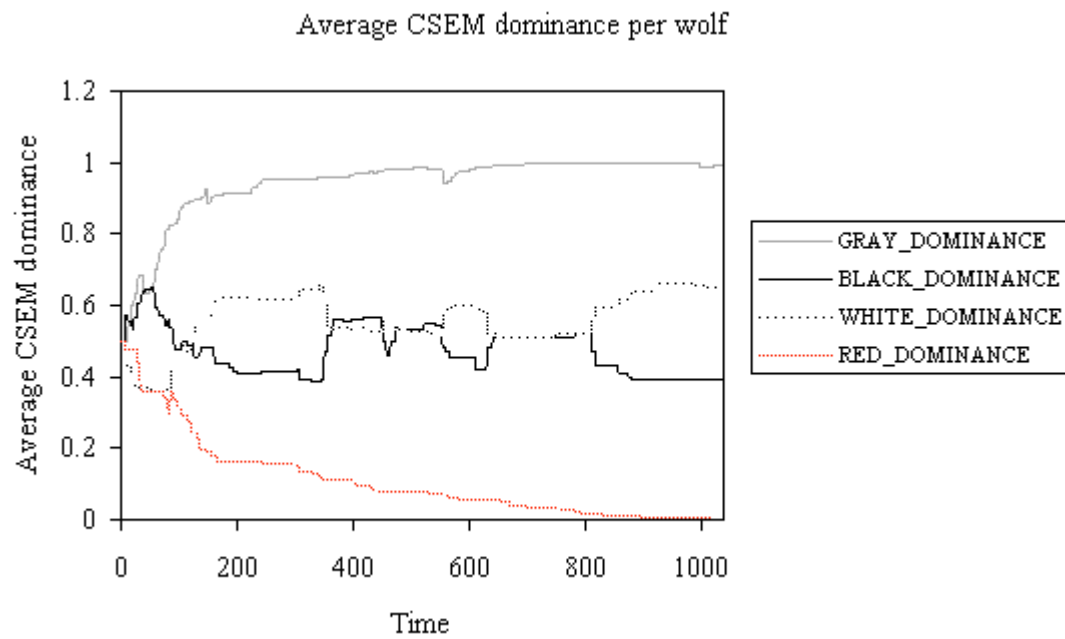
Figure 4-12 shows a group of six virtual wolves who started off at  $t = 0$  with no relationships towards each other. A small amount of initial randomness caused by arbitrarily assigned initial positions becomes magnified due to the positive feedback system of the synthetic relationships. The blue pup rapidly emerged as the most dominant, dominating all other pups by  $t = 500$ . Gray was not far behind, dominating everyone but blue on the same time scale. The four lower ranking wolves took a bit longer to sort out their relationships but eventually (by  $t = 5000$ ) settled into a fairly stable configuration. Green dominated everyone but blue and gray; black submitted to blue, gray and green; white submitted to everyone but red; red, at the bottom of this emergent hierarchy, submitted to all the other wolves. This graph will be discussed further in the Evaluation chapter (section 5.3.1).



**Figure 4-12: A continuous representation of the pups' dominances.**

## 4.8.2 An Example

As an example of how CSEMs and their visualizers work, consider the following graph of the average dominances among a four-member pack (see Figure 4-13).




**Figure 4-13: The first 1000 virtual seconds in a four-member pack.**

The four pups – “gray”, “black”, “white” and “red” – all start off life identically, with no relationships towards each other. As the pups meet each other, they begin to form relationships. Gray rapidly establishes itself as a dominant individual, and red emerges as the most submissive. White and black haggle over the middle ranks for the first 800 virtual seconds, and then settle into a somewhat more stable configuration.

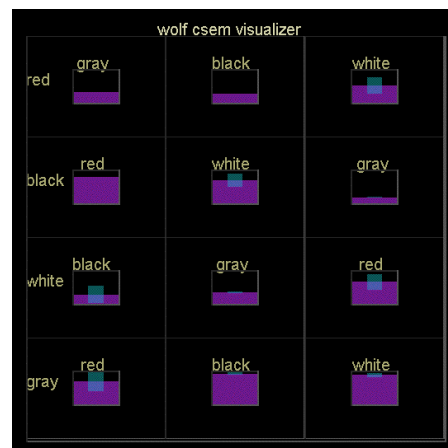
The next few figures show the instantaneous CSEM visualizers at various points in the same four-member litter of pups.

Figure 4-14 shows their relationships at approximately  $t = 50$ . All pairs of pups have met each other except white and red (visible because neither pup has a CSEM for the other). None of the pups are very confident in any of their relationships, though, since they’ve just been formed – all of the CSEMs show a lot of blue.



**Figure 4-14:  $t = 50$ . The pups are just getting to know each other.**


By  $t = 100$  (see Figure 4-15), all the pups have met each other. Gray is beginning to show signs of dominance within the group; black has established itself over red; all of the other relationships are still in a state of flux.



**Figure 4-15:  $t = 100$ . All the pups have met each other. A few of the relationships have settled in.**

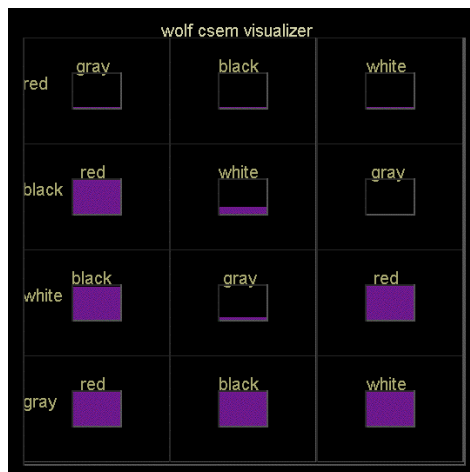
By  $t = 400$  (see Figure 4-16), all of the relationships are fairly well-established except the one between white and black. While the relationships are not fully polarized (with one individual

exhibiting dominance  $\approx 1.0$ , and the other with dominance  $\approx 0.0$ ), they nevertheless show clear and agreed-upon separations between the dominance roles of the two pups.



**Figure 4-16:  $t = 400$ . All of the relationships have stabilized except the one between white and black.**

Figure 4-17 shows the fully resolved relationships among the pack. Gray is clearly the “top dog”, black has conceded to white, and red is solidly submissive to all comers. As I demonstrate in section 5.3.1, reversals can happen (especially among middle-ranking wolves), but on the whole, the relationships will stay fairly stable once they are established and all individuals are confident in them.



**Figure 4-17:  $t = 1000$ . The dominance relationships are now well established, with gray at the top, then white, then black, and red at the bottom.**

### 4.8.3 The Three Implementations

There were subtle differences among the CSEM implementations in the three projects that featured the AlphaWolf system. The AlphaWolf installation featured the implementation described here. The user tests and resource exploitation simulations contained extensions to allow them to run each of the four social relationship algorithms described in section 5.2.2.1, and the apparatus to read CSEM data from a file and write it to a file. Nevertheless, the AlphaWolf mechanism itself remained essentially the same in the user test implementation. The simulations had a stronger interaction model, in which wolves instantaneously applied their CSEM (i.e.,

regardless of proximity), and revised their CSEMs as soon as they ceased reciprocal attention (i.e., whenever one of them stopped paying attention to the other.)

## **4.9 Development**

The wolves in the AlphaWolf installation grew up from pups to adults over the course of each five-minute installation. Their development had two main components – physical and behavioral.

### **4.9.1 Physical Development**

The pups changed shape over time due to a morphology system written by Scott Eaton [Eaton 2001] and adapted by Marc Downie. This system generated the wolves' geometries and animations in real-time by blending two sets of models and animations, one of a pup and the other of an adult. A straightforward blend of the animations, while adequate, did not always look good enough. For example, because of differences in the proportions of the pups' and adults' legs, adolescent wolves' feet sometimes went through the ground during a blended walk. To remedy this problem, Downie wrote a system that prevents the feet from going below the level of the ground. Nevertheless, the problem of developmentally blended characters and animations continues to be a difficult problem.

In addition to the model and animations reflecting developmental changes, wolves' voices changed as they aged, too. Pups start off with high pitched howls, growls and other sounds, and gradually mature into deeper, more resonant vocalizations. The sounds made by the wolves will be discussed further in section 4.10.2 below.

### **4.9.2 Behavioral Development**

As well as changing shape over time, wolves also changed their behavioral patterns as they grew up. For example, certain adults would go off hunting for a portion of each run, while the pups always stayed near the center of the virtual forest (the “den”). The system that controlled this behavioral development was quite *ad hoc*, involving an assortment of age-based TriggerContexts for each ActionTuple. A possible formalization of this behavioral development mechanism is presented in the Future Work chapter (section 7.2.10).

While it was not explicit, the increasing confidence that wolves tended to show in their CSEMs also served as a mechanism by which wolves matured. This social maturation was loosely matched to the learning curve of the novice users who were controlling them. As the users became more comfortable with the interaction paradigm, the pups became more confident in their relationships.

### **4.9.3 Visualizer**

The bodies and behaviors of the wolves themselves served to visualize their developmental changes, so no additional visualization tools were necessary.

### **4.9.4 The Three Implementations**

The AlphaWolf installation at SIGGRAPH featured the physical and behavioral development described above. The user tests and simulations had no adults, and therefore no developmental component. However, all three had the emergent social maturation that occurred through the gradual reduction of the CSEMs' confidences as the pups worked out their relationships.

## 4.10 Expression

One of the core elements of our system is the strong graphical expression of the behaviors of the wolves. Rather than being numbers on a screen or even dots on a flat plane, the wolves are fully rendered, 3-D animated, sound-producing virtual characters. The two main parts of the expression system of the wolves are the animation system and the sound system.



**Figure 4-18: An emotional range from submissive to dominant. Our animator crafted the extreme poses. The center three poses are automatically generated.**

### 4.10.1 Animation

The expressive range of animation for the characters is essentially the “front end” for the emotion system described in section 4.6 above. Each wolf’s emotional state affects what it does and how it does it. Our system uses the work of Marc Downie [Downie 2001a] to let our characters have dynamic expressive ranges (see Figure 4-18). Downie’s expressive motor system allows a wolf’s current emotional state to affect the style in which the wolf behaves (e.g., run dominantly or submissively). Since emotional memories affect the wolf’s current emotional state, they also affect the style of the wolf’s behavior. While the motor system is not the focus of this thesis, it is an important element of the social relationship system in the wolves. The most sublimely complex representation of a social relationship isn’t of much use without an equally expressive range of behavior.

The source animations of the wolves were created in 3D Studio Max by our talented animator Adolph Wong. The main reference material used in creating the animation were a number of videos depicting wolves in the wild (e.g., [Rosenfield 1988; IMAX 1999]), as well as pictures from various books (e.g., [Busch 1998; Coppard 1999]) and images found online. Drawing inspiration from nature in the animation, as well as in the entire system design, helped keep the project true to its natural model.

The animations were exported from 3D Studio Max using a custom exporter written by Matt Berlin. This exporter converted the animations into a format readable by our system.

In terms of animation, the central expression of dominance and submission comes from the spinal column. As Figure 4-18 shows, the arching of the back and holding upright of the head account for much of the dominant look of the right-hand pictures. In addition, the ears and tail contribute significantly to the emotional mood of an animation. Finally, the motion itself, as well as the poses, captures much emotion; for footage of the wolves’ motion, please see the online video of AlphaWolf at the following address:

<http://badger.www.media.mit.edu/people/badger/alphaWolf/alphaWolf.mov>

The source animations involved cycles (e.g., “walk”, “sit”), transitions (e.g., “sit\_to\_walk”), and layers (e.g., “face\_growl”). Most of these animations have versions for dominant pups, submissive pups, dominant adults and submissive adults. By blending along these two axes (age



and dominance), a full range of emotion and development were simultaneously generated. In addition, certain animations (e.g., “walk”, “gallop”) had directional versions as well (i.e., “left”, “straight”, “right”). Thus, our most complex blending occurred among 12 examples (2 ages x 2 emotions x 3 directions).

#### 4.10.2 Sound

Certain animations created sounds when they were played. These sounds were sent as MIDI notes to a Korg synthesizer, in which were stored all the various wolf samples. For example, the suite of howl animations (e.g., “pup\_howl\_standing\_dominant”) made a howl sound whenever a pup played one of them. Often this sound was a randomized choice of one of several sounds, played at some slight variation of its original pitch. As mentioned above, the sounds changed developmentally as well.

Each pup’s sounds were played through one of the installation’s four speakers. This localization helped build a relationship between the participant and his or her wolf, because that wolf’s sounds came from the speaker nearest to that participant. In addition, it helped build the relationship between different people interacting, since participants could localize the pups’ noises in the real world as well as the noises the participants were making into the microphone.

Most of the wolves’ sounds were derived from real wolf vocalizations, primarily from a collection of sounds purchased from the Macaulay Library of Natural Sounds at the Cornell University Laboratory of Ornithology.

The sounds made by the adults, and the ambient wind noise of the virtual world, were distributed among the four speakers. Individual and diverse gusts of wind played from specific speakers, giving an immersing ambient sound-scape supporting the wolves’ vocalizations.

A sound experiment that did not make it into the final version of AlphaWolf was a barely audible, bass-heavy heartbeat that increased its frequency in response to social stress in the pack. While it did appear to strengthen the emotional impact of the installation, some people found it a bit creepy, so it was removed for the final version.

The system that generated the wolves’ sounds was created by Bryan Yong and Marc Downie. Bryan Yong performed the manipulation of the samples that created the emotional and developmental ranges.

#### 4.10.3 Visualizer

The images and sounds of the installation are the visualizers for the expressiveness component of the implementation.

#### 4.10.4 The Three Implementations

Each of the three implementation drew on the same body of animations and sounds. However, for the user tests, the connection between the emotional state and the expression of that emotion was made more extreme. In the user tests, the following function was applied to the Dominance emotion before it was fed into the motor system:

$$\begin{aligned} \text{if } d < 0.5, d' &= ((-1 * \sqrt{|d - 0.5| * 2}) + 1) / 2 \\ \text{else, } d' &= (\sqrt{|d - 0.5| * 2} + 1) / 2 \end{aligned}$$

Equation 4-5

where  $d'$  is the dominance value used by the motor system, and  $d$  is the dominance value used by the behavior system. This function effectively increased the contrast on the expression of the wolves' emotional ranges in the user tests. The simulations shared this more extreme expressiveness.

## 4.11 Cinematography

The wolves were displayed on AlphaWolf's four screens through an automatic cinematography system based on the author's previous work [Tomlinson 2000b]. Cinematography is a social medium – camera placement helps to define the relationships among actors. To show off the social nature of the characters, the camera system attempted to frame shots around two interacting characters. In addition to being a showcase for social characters, AlphaWolf was an interactive experience; therefore, the camera system needed to facilitate each participant's ability to see where his or her pup was going.

The camera system attempted to blend the two motivations above into a seamless experience for participants. Many computer games (e.g., *The Legend of Zelda*) have different “modes” for the different kinds of interactions that a character might have – one camera style for talking to people, another for fighting things, a third for navigation. I have found the use of modes to be disconcerting and to break the suspension of disbelief, especially when it is left to the player to choose the mode. For AlphaWolf, I wanted an unobtrusive camera system that would nevertheless make it easy to navigate and cleanly display the relationships that the “leading pup” formed.

The essential camera plan was:

- 1) Pick a shot (camera position and camera target).
- 2) Hold it for either two seconds or until it ceased to adequately capture the moment, whichever was longer.
- 3) Pick a new shot, and smoothly transition to it.

Criteria for “until it ceased to adequately capture the moment” included:

- 1) the lead pup is no longer on screen.
- 2) the lead pup is very far away from or very close to the camera.
- 3) the pup is facing the camera for more than a second.

In addition to this general plan, the camera performed some “exception handling.” For example, if the camera's preferred position placed it within the geometry of a pup, the camera smoothly (but rapidly) moved up until it was above that pup. While this move sometimes reduced the quality of the shot, it avoided the awful occurrence of seeing through a camera located inside of a computer graphical object. Few things in computer graphics are more jarring than seeing the inside of a 3D-animated character, especially in a piece meant to be fluid and evocative.

A second case in which the camera left the central plan was when a wolf was directed by its participant to leave the center of the forest (a circle with radius of approximately 20 wolf-body-lengths). This central area was large enough for all the wolves to interact in very comfortably, and there was nothing to see or interact with outside of that perimeter (though the world did continue to “exist”, i.e., there was no black void at the edge of the universe). If a wolf left this central area, the camera would arrange itself to look at that wolf in the direction of the center of the world (see Figure 4-19). This placement had the effect of causing the vast majority of screen-space on which the participant might click to direct the wolf back towards the center of the world.

While the participant could force the pup further and further out by carefully clicking right in front of the wolf, it rapidly became clear that there was to be no reward for such behavior.



**Figure 4-19: The camera angle that guides people back to the center of the world.**

Each of the three kiosk-cameras worked with the above plan; the central plasma display camera had a variant on it. The plasma camera used the same plan but chose whichever lead actor appeared to be engaged in the most “interesting” interaction, where interesting is defined as “nearest to another individual.” This choosing of the lead actor was anti-aliased to prevent the camera from switching rapidly between pups. The plasma camera also stayed a bit further away from the action, so that the pups were not enormous on the 52” display.

Wolves also had the ability to look at the camera on cue. In particular, they do so right after waking up. This direct connection between wolf and player early on in the interaction appeared to be very helpful in building the human-computer relationship.

#### **4.11.1 Visualizer**

Because the camera system is its own visualizer, there was no need to have a visualizer for it. However, during previous synthetic camera system projects [Tomlinson 2000b], the “visualizer” for the camera system was a set of sound effects that played when the camera was doing a certain behavior (e.g., “Action Shot”, “Emotion Shot”) so that the developer did not need to be watching a text output system while trying to examine the visual motion of the camera. The fact that most of the visualizers in the AlphaWolf system are purely visual, rather than engaging the other senses, could be an area for improvement.

#### **4.11.2 The Three Implementations**

The full AlphaWolf installation had four screens with the two camera paradigms described above. The user tests, designed for only one participant with no audience, had only a single screen, which employed the kiosk-camera paradigm and focused on the gray wolf pup. The simulations had a similar camera/screen set up to the user tests, although no one was watching during their runs.

## **4.12 Interface**

People direct the AlphaWolves through two main interfaces – a microphone and a mouse, which are discussed below.

### **4.12.1 Microphone**

For an installation that showcases social computational entities, it seemed appropriate to have an interface that was closely connected with sociality. For AlphaWolf, we chose a vocal interface (the “Howl Interface”) since the voice is one of the most evocative communication mechanisms at a person’s disposal. At very least, we thought it would be better than using a keyboard. (In the user study, we test how the Howl Interface compares to using buttons for the same tasks. See section 5.2.3.3.)

Each participant’s microphone is backed by a system running on a separate computer that performs acoustic pattern matching on the utterances that it receives. We use a simple mechanism for classifying sounds involving utterance length and harmonicity. Howls are long and harmonic. Whines are short and harmonic. Growls are long and non-harmonic. Barks are short and non-harmonic. A fifth, “silence” classifier kicks in if the volume coming into the microphone is below a certain threshold.

The Howl Interface uses simple short time Fourier transform windows to classify the audio into one of four categories based on the length of an vocalization and an estimation of its harmonicity. Harmonicity is determined by looking at the percentage of energy in a window around a harmonic of the most prominent frequency. The threshold for long utterances is 0.5 seconds. In addition to simply classifying the sound, the volume of the sound was also sent to the system; louder growling, for example, provoked more vigorous dominant behaviors from the pup. This simple and elegant system, written by Jennie Cochran, 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 show floor.

### **4.12.2 Mouse**

When a user clicks on the screen with the mouse, the wolf moves toward that point. Matt Berlin wrote the system that casts a ray from the camera’s position to the point clicked, and determines where that ray intersects with the ground plane. If the intersection point is within a certain distance (approximately 10 wolf-body-lengths), the wolf will run to the point clicked and sniff the ground there. If the ray intersects the ground further than that distance away, or if the click does not intersect with anything (i.e., the click is above the ground plane), the pup will run that distance in the direction of the click and sniff the ground.

If the participant clicks on another wolf, that wolf becomes the target to which the pup runs.

If the participant clicks on one of the buttons at the top or bottom of the screen, the pup runs to the place where it last remembers having encountered that social partner. If it sees the social partner on its way to that spot, it will recognize that wolf and run to it. Otherwise, it will run to the place where it *expected* it to be and sniff that spot.

### **4.12.3 Visualizer**

For the Howl Interface, there are three visualizers, one for each pup/microphone combo. Each visualizer is a window with a column of words in it, rapidly updating. The words are either

“None”, “Howl”, “Growl”, “Whine”, or “Bark.” Because the system updates rapidly, it is easy to see what classification is taking place and the duration of each utterance.

#### **4.12.4 The Three Implementations**

The above system was used in the AlphaWolf installation, and as one of the portions of the user tests. In addition, the user tests had a system that could replace the microphone with two keys on a keyboard which, when pressed, served the same function as growling and whining. The simulations had no interface for providing directorial input except for the mechanism that allowed a human or other entity to force a certain wolf to act consistently dominant.

#### ***4.13 Summary***

To summarize the essential elements of the implementation of the virtual wolves’ social relationship mechanism:

- Each virtual wolf is an entity.
- When another wolf enters its perceptual environment, it forms a Context Specific Emotional Memory (CSEM) of the other wolf (now a “social partner”).
- Interactions with the social partner affect the wolf’s emotional state.
- At the end of each interaction episode, the CSEM is revised based on the wolf’s emotional state at that moment.
- On successive encounters, the CSEM is applied, affecting the wolf’s emotional state in proportion to its confidence in the CSEM.
- The emotional state affects both the wolf’s choice of actions (if it is not user-controlled) and the expressive style in which the wolf takes its actions.
- Each wolf grows up over time, both physically and behaviorally.
- The human interface allows a person to control the actions of the wolf, and thereby to direct the relationships that the wolf forms with its pack mates.
- Supporting technologies including camera and sound help enhance the believability and directability of the virtual wolves.

## 5 Evaluation

This chapter presents the three ways in which we evaluated the social relationship mechanism and the rest of the AlphaWolf system. These three parts are: a set of human user studies, a set of computational simulations, and the reception of the AlphaWolf installation at SIGGRAPH. Section 5.1 summarizes the results of all three of these parts. Section 5.2 describes the human user studies in full. Section 5.3 details the set of simulations that demonstrate the effectiveness of the AlphaWolf mechanism in the domain of social resource allocation. Section 5.4 offers some thoughts from the run of the AlphaWolf installation at SIGGRAPH.

### 5.1 Summary of Results

In this section, I summarize the clearest and most important results achieved by the various parts of this evaluation.

#### 5.1.1 Transmission of Relationships

**Hypothesis:** The AlphaWolf mechanism will encode social relationships created by a first person in a way that can be perceived by a second, naïve person.

In the user tests, each of 32 subjects directed the gray pup to form specific social relationships with its siblings, the white and black pups. The CSEMs of each pup were recorded at the end of the run. The CSEM representation matched the relationship that had been assigned to the user in 93.8% of the cases ( $p < 0.0001$ , chance would be 50%), demonstrating that people were successful in forming relationships as directed. These same CSEMs were then loaded into a different litter of pups at the beginning of a later subject's run. These second subjects succeeded in recognizing 86.7% of the relationships that they viewed ( $p < 0.0001$ ), demonstrating that the wolves were clearly expressing their relationships. The AlphaWolf system succeeded in communicating the relationships from the first subject to the second subject in 89.7% of the cases ( $p < 0.0001$ ).<sup>3</sup>

**Result:** The AlphaWolf mechanism successfully encodes, decodes and transmits a valid representation of social relationships.

This result suggests that the CSEM representation described in this thesis is a simple, usable representation of a social relationship. People could direct their pups to form relationships; therefore, the AlphaWolf system is amenable to human control. People could perceive the relationships; therefore, the AlphaWolf system is clear and expressive. While this test is very simple, and only a first step towards a full model of a social relationship, the results confirm that it is a solid start.

#### 5.1.2 Effect of Social Relationship Algorithms

**Hypothesis:** The AlphaWolf algorithm will be preferred by users to alternate social relationship algorithms in a variety of areas of subjective experience.

---

<sup>3</sup> The percentage transmitted was higher than the percentage decoded due to a double error, in which an incorrectly encoded relationship was in turn incorrectly decoded, resulting in successful transmission. Since this run was discarded, the denominator was decreased by one, resulting in a higher percentage.

Results in this second area tested by the human-user studies were apparently swamped by the effect of the order of runs. Fully 81.9% of subjects preferred the second of two runs where the only difference was the social relationship algorithm ( $p < 0.0001$ ). Nevertheless, there was some evidence that AlphaWolf was preferred over the three alternate algorithms (see section 5.2.2.1 for a full description of the alternate algorithms), with 67.9% of the responses that favored the first of the two runs choosing the run with the AlphaWolf mechanism ( $p = 0.0436$ ).

**Result:** There is strong evidence that the effect of the algorithm on users' subjective experience was less powerful than the effect of novelty. There was weak evidence in support of AlphaWolf creating a superior subjective experience for participants.

It is interesting that people didn't have a strong preference with regard to algorithm, when the algorithm clearly had a strong impact on the results pertaining to the transmission of relationships. I believe that the effect of the various social relationship algorithms would become more apparent in a longer interaction or with experienced users, where the effect of novelty was not as substantial.

### 5.1.3 Effect of Interaction Paradigms

**Hypothesis:** Interactive versions of AlphaWolf will outperform non-interactive versions in a variety of capacities.

This is the third and final topic addressed by the human-user studies. There was strong support for various components of this hypothesis and strong contradictions of other aspects. In particular, interactivity (both Button-controlled and Microphone-controlled versions) allows subjects to feel substantially more control and creates a somewhat more immersing experience. Nevertheless, interactivity appears to reduce the realism of the virtual wolves, and to decrease subjects' feeling that the relationships among the pups were clear.

A significant related result is that people appear to like a wolf more and identify more closely with it when that wolf is the camera's object of attention. The effect of the camera is greater than that of interactivity; people like and identify with the gray wolf only slightly more when they can direct it than they do when they're just watching. Within the interactive runs, the Microphone interaction paradigm appears to increase the identification that people feel with the pups over the level found in the Button runs, although it doesn't make them like the pups more.

**Result:** Interactivity and cinematography both make a significant difference in a variety of capacities to participants in a novel experience, though not always a positive difference.

The interaction paradigm, unlike the social relationship algorithm, had a significant impact on users' subjective experiences of a completely novel interaction. In designing interactive installations and other applications, the way people interact with it could make more of a difference than what's going on "under the hood," at least at first glance. This result is unsurprising to us, since one of the things that really sticks in people's minds about the AlphaWolf installation at SIGGRAPH is the Howl Interface, despite the presence of many other equally novel and complex elements in that installation.

For a full treatment of each of the three topics above, please see section 5.2 below.

#### 5.1.4 Simulations of Resource Exploitation

**Hypothesis:** The AlphaWolf mechanism is a simple, robust, social relationship mechanism that can be used as the basis for relationships among entities in a multi-agent or multi-robot system.

In addition to the user tests, I performed a series of small experiments in simulation, to ascertain the usability of the AlphaWolf mechanism from the point of view of a creator of multi-agent systems. Specifically, I created a pack of 6 virtual wolves to examine their effectiveness in resource exploitation. These simulations demonstrated that the AlphaWolf mechanism creates dynamic hierarchical social structures that are similar to those of real wolves, and outperforms three alternate mechanisms in this regard. In addition, they demonstrate that groups of virtual wolves with the AlphaWolf mechanism exhibit a reliable disparity in resource allocation among its members, which might be usefully applied to multi-agent or multi-robot systems. Finally, they demonstrate that the multi-agent systems created with the AlphaWolf mechanism are directable by an external entity, also a useful characteristic for multi-agent and multi-robot systems.

**Result:** The AlphaWolf mechanism is an effective basis for a range of simulations of hierarchy-formation, resource-exploitation, and directable multi-agent systems.

While the simulations implemented for this section represent only a tiny subset of the range of wolf social behavior, the AlphaWolf system's effectiveness in accurately capturing this subset gives hope that the representations and ideas described in this thesis will continue to be of use in a wide range of other possible applications.

For a full treatment of these simulations, please see section 5.3.

#### 5.1.5 The AlphaWolf Installation at SIGGRAPH

**Hypothesis:** Visitors to SIGGRAPH will enjoy the AlphaWolf installation, in part because of the social relationships that are formed there.

The third way in which we judged the AlphaWolf system was to take it to SIGGRAPH and put it to the test in front of a large audience. Over the course of 5 days, between 500 and 1000 visitors participated in the AlphaWolf installation.

**Result:** Visitors to SIGGRAPH did appear to enjoy the installation, and participated vigorously in the wolves' social relationships.

While we did not collect any data while at SIGGRAPH, we observed people's interactions with the installation and with each other. For a longer description of SIGGRAPH visitors' engagement with AlphaWolf, please see section 5.4.

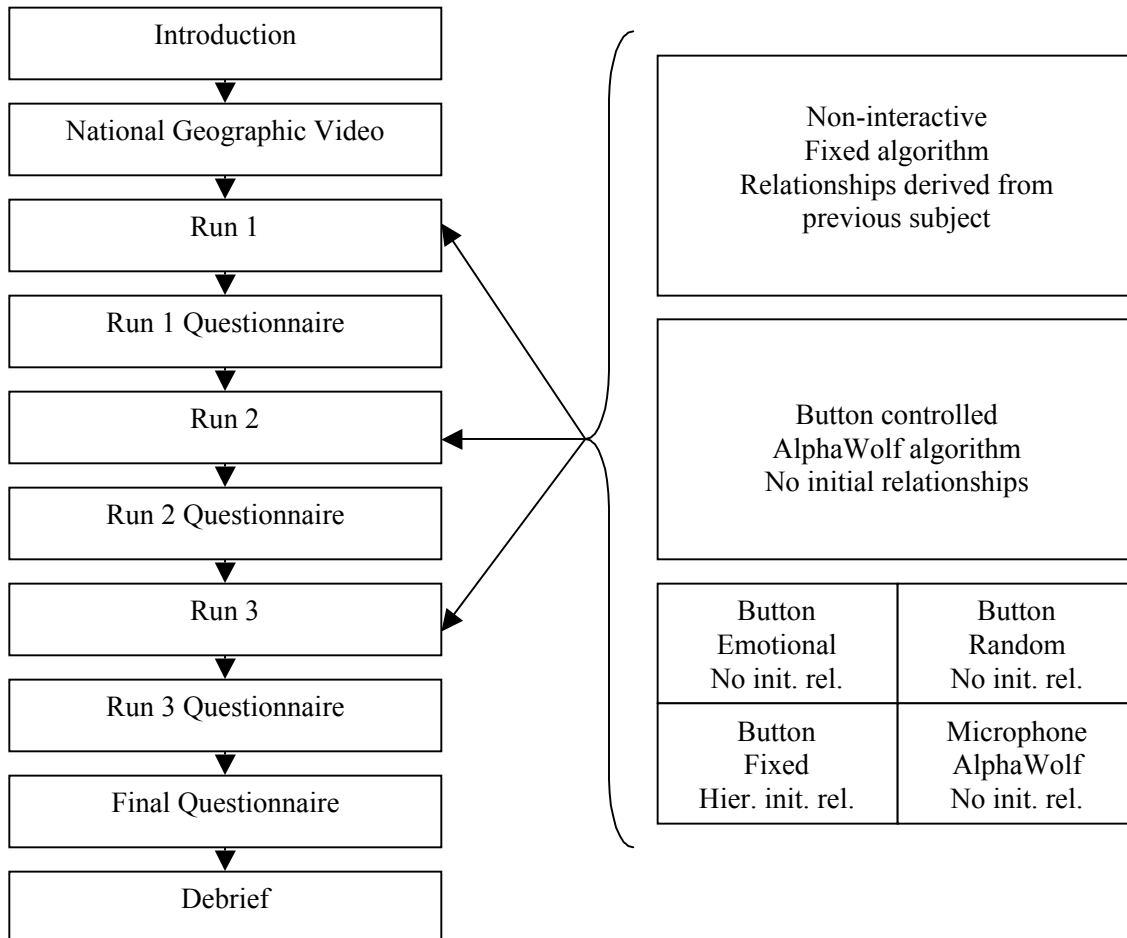
### 5.2 User Tests

In order to examine the effectiveness of the AlphaWolf system and its social relationship mechanism, I conducted a series of human-subject experiments. In these tests, users watched and interacted with a selection of controlled runs of the virtual wolf pack. Sections 5.1.1, 5.1.2, and 5.1.3 present summaries of the three main areas examined in these user tests.

As a broad overview, the study involved 32 subjects, each of whom came to our Lab for a 45 minute period, watched a short clip from a National Geographic Video, and interacted with several runs of virtual wolves (see Figure 5-1). Each run featured three wolf pups – one gray, one



black and one white. Each run was driven by one of four different social relationship algorithms (described below), and included one of three interaction paradigms (described below). After each run, the user was asked to fill out a questionnaire ranking his or her opinion on a range of topics. At the end of each session, the subject was asked to fill out one final questionnaire comparing the various runs. This experiment was approved by the MIT Committee on the Use of Humans as Experimental Subjects (COUHES # 2864). In addition to the main study, I performed a pilot study with 9 subjects to work out several issues in the experimental design.



**Figure 5-1: Diagram of human user study's experimental method.**

The user tests were designed to evaluate **three main research areas**:

- Does the AlphaWolf system capture some essence of social behavior (i.e., does it work)? In particular, is the CSEM an effective computational representation of a social relationship?
- How does the AlphaWolf algorithm compare to other possible social behavior algorithms? What effect do these algorithms have on a participant's subjective experience?
- How much of an impact does the interaction paradigm have on a user's experience of the wolves? What effect does the interface have on a participant's subjective experience?

### 5.2.1 Hypotheses

In order to establish concrete, falsifiable hypotheses for the study, I distilled the three research areas above into three main groups of hypotheses:

**Hypothesis Group 1:** The AlphaWolf mechanism encodes social relationships created by a first person in a way that can be perceived by a second, naïve person. This hypothesis is, in fact, a combination of two sub-hypotheses: a) The AlphaWolf mechanism successfully encodes social relationships, and b) Naïve users can effectively perceive the relationships in our system.

**Hypothesis Group 2:** A run with the AlphaWolf algorithm will outperform a run with the Emotional algorithm, which will in turn outperform a run with the Fixed algorithm, which will in turn outperform a run with the Random algorithm (i.e. AlphaWolf>Emotional>Fixed>Random), in each of the following areas: similarity to real wolves, clarity of social relationships, user enjoyment, user control, user comfort with the interface, user liking of the wolves, user identification with the wolves, and user immersion.

**Hypothesis Group 3:** A run with a microphone interface will outperform a run with a button interface, which will in turn outperform a non-interactive run (i.e., Microphone>Button>Non-interactive), in each of the following areas: similarity to real wolves, clarity of social relationships, user enjoyment, user control, user comfort with the interface, user liking of the wolves, user identification with the wolves, and user immersion.

### 5.2.2 Experimental Method

In this section, I describe the alternate algorithms and interaction paradigms, explain the procedure that was followed with each subject, describe the group of subjects, and depict the virtual wolves with whom those subjects interacted.

#### 5.2.2.1 Alternate Algorithms

In addition to the AlphaWolf algorithm which is described in depth in the Implementation chapter, I implemented three other mechanisms of social relationship formation, which I will call Emotional, Fixed, and Random.

##### Emotional

The Emotional algorithm is identical to the AlphaWolf mechanism except that the emotional memories are prevented from influencing the wolf's current emotional state. Referring back to the Implementation->Learning section above, the "apply" function of each CSEM was disabled. All else was held the same – memories were formed identically, emotions influenced action selection. The observed effect of this algorithm is that the pups maintained their current emotional state until they had an interaction with another wolf that changed it. Because interactions often end with one wolf at each end of the dominance range (one dominant and one submissive), the Emotional Algorithm caused the wolves to spend most of their time at one of the extreme (and therefore hand-animated) emotional states.

##### Fixed

The Fixed algorithm loads in a predetermined set of CSEMs that specify the relationships among the three different pups. These randomly assigned static relationships are similar to Goldberg's algorithm in which individuals win dominance interactions based on a pre-assigned unique ID [Goldberg 1997].

The predetermined social relationships took one of two forms.

1) If the run was a Non-interactive run (see below), the relationships were loaded from a file. This file contained a representation of a previous pack's set of relationships created at the end of some previous subject's interactive run. This two-stage process allowed us to test whether the subject viewing that Non-interactive run could detect the relationships that had been encoded by that previous subject during his or her interactive run.

2) If the run was a Button run (see below), the relationships were loaded from a file that encoded a pre-specified linear dominance hierarchy. The relationships in this case were very strong (Value = 0.0 or 1.0, Confidence = 1.0), and were matched such that each wolf was either completely dominant or completely submissive to another (e.g., no pair of wolves both think they're dominant or both think they're submissive).

In the Fixed algorithm, CSEMs are applied just as they are in the AlphaWolf algorithm, but they are never revised (again, see Implementation->Learning). They stay exactly the same as they were when they were loaded, regardless of the interactions that occur between each pair of wolves.

### **Random**

The Random algorithm picks a random value between 0.0 and 1.0 for each CSEM every time the revise function of that CSEM is called. Confidence on all CSEMs is held at 1.0. The effect of this algorithm is that every time two wolves meet, they have a completely new relationship. In addition, there is no verification that a pair's relationships "match." It is perfectly possible that two wolves could both feel completely dominant or both completely submissive toward each other.

### **5.2.2.2 Interaction Paradigms**

The three different kinds of interaction paradigm that people experienced in the experiment are: Microphone, Button, and Non-interactive.

#### **Microphone**

The Microphone interaction paradigm for the user tests was a stripped-down version of the microphone interface for the AlphaWolf installation. The utterance classifying system in AlphaWolf detected four kinds of sound – howl, growl, whine and bark – as well as silence. In the user tests, the interaction was kept as simple as possible, so that there would be no confusion about directions. Therefore, I revised the system to recognize only growls and whines. Howls were categorized as silence, and barks counted as growls. Aside from this modification, the system was identical to the microphone input used at SIGGRAPH. In addition to the microphone itself, users were given a mouse with which to tell their pups where to go (just as in the full AlphaWolf installation).

#### **Button**

The Button interaction paradigm replaced the two vocalizations enabled by the Microphone – growl and whine – with buttons on the keyboard. Two keys were labeled with signs reading "GROWL" and "WHINE"; each caused the user's pup to take the corresponding action. As in the Microphone case, users were given a mouse with which to tell their pups where to go.

#### **Non-interactive**

Users are asked simply to sit and watch the wolves, and do not interact in any way.

### 5.2.2.3 Procedure

This section describes the entire process of a single subject participating in the study. (see the Appendix for full experimental method). All experiments took place in the Synthetic Characters' lab space in the MIT Media Lab's One Cambridge Center office complex. Figure 5-2 provides an image of the experimental set-up.



**Figure 5-2: The experimental setup.**

### Preparation

Before a subject arrived for a session, the experimenter prepared the physical space and readied the various stages of the experiment, so that all subjects would have as uniform an experience as possible. This process involved:

- setting up the batch files that would launch each run with a single click.
- verifying that the computer screen was clear of any extraneous windows, and that the surround sound system was turned on and set to the appropriate volume.
- preparing the “table tents” that would be given to the subject to remind him or her of the controls and relationships that the pup should form.
- numbering the various forms and placing them on the table where the subject would be sitting.
- verifying that the room was clean and the curtains were closed.
- putting the “Experiment in Progress” sign on the door.

### Introduction

When the subject arrived, he or she was greeted and invited to sit down in front of a computer screen. The room featured a variety of computers and wolf paraphernalia (drawings, posters, etc.), but was otherwise tidy. The desk where the subjects sat had the following items on it: a keyboard, a mouse, a microphone, a stack of questionnaires with the consent form on top, a folder for putting completed questionnaires in, a notepad, several pens, and the sleeve to a National Geographic Video they would later be shown (see below). Aside from the elements described below, all aspects of the experimental space (e.g., thermostat, etc.) were held constant.

At the beginning of the session, the subject was read a brief description of the various stages of the experiment, as well as the conditions of the experiment (“You are welcome to stop at any time”, “You will receive a \$10 gift certificate to Toscanini’s Ice Cream for your participation”, etc.) The subject was then asked to read and sign a consent form (see Appendix). Thereafter, the subject was shown the questionnaire that he or she would be asked to fill out later, to give him or her an idea of what factors to attend to in the virtual wolves.

### **Video of Real Wolves**

Every subject was first shown a short, 2.5 minute video of real wolves in the wild. This video was a selection of clips edited together from the National Geographic Video “White Wolf” [Rosenfield 1988]. This video showed a pack of arctic wolves (including a litter of pups) engaging in dominance interactions, and featured narration by noted wolf expert David Mech. After starting the video, the experimenter left the room, and instructed the subject to knock on the door when the video finished.

### **The Three Runs**

After the subject knocked, the experimenter re-entered the room and read the directions for the first of three runs. Each of the three runs featured identical packs of virtual wolves except for several parameters, described below, that were randomized.

1. Every subject had one Non-interactive run, featuring the Fixed algorithm. This run had pre-loaded social relationships that were the result of some previous subject’s Button run.
2. Every subject had one run where he or she directed the gray pup with the buttons. All three wolves in this run featured the AlphaWolf social relationship mechanism. This run will be called “Button/AlphaWolf” throughout the rest of this document. In the Button/AlphaWolf run, the relationships among the various pups started off identically, with each pup having a CSEM with dominance = 0.5 and confidence = 0.5 for each other pup.
3. Each subject had a third run, randomly selected from the following four combinations:
  - a. “Button/Emotional” – A button interaction run in which the wolves feature the Emotional social relationship mechanism. Holding the interaction paradigm constant (with regard to run 2, above) and varying the algorithm served to isolate the effect of the algorithm. This kind of run started off with each pup having an emotional memory of 0.5 dominance and 0.5 confidence towards each other pup.
  - b. “Button/Fixed” - A button interaction run in which the wolves feature the Fixed social relationship mechanism. This kind of run started off with each pair of pups having a strongly polarized relationship (i.e., one pup had a remembered dominance of 1.0 and confidence of 1.0, the other had dominance 0.0 and confidence 1.0). The relationships were arranged such that the pups formed a linear dominance hierarchy ( $A > B > C$ ), though the order of the pups in this hierarchy was randomized. Through the Fixed algorithm, these relationships remain unchanged.
  - c. “Button/Random” - A button interaction run in which the wolves feature the Random social relationship mechanism. This kind of run started off with each pup having an emotional memory of random dominance (between 0.0 and 1.0) and confidence of 1.0 towards each other pup. The Random algorithm caused these relationships to change frequently.

- d. “Microphone/AlphaWolf” – Holding the algorithm constant and varying just the interaction paradigm served to isolate the effect of the interface on the subject’s experience. This kind of run started off with each pup having an emotional memory of 0.5 dominance and 0.5 confidence towards each other pup.

As the above descriptions show, two of the three runs were interactive; in these runs, the subject played the role of a gray wolf pup. In addition to the gray pup, there were two other pups, one black and one white, controlled by autonomous behavior systems. The gray pup was chosen as the pup the subjects control because it seemed to minimize the effect of color on the various elements we hoped to measure. In Non-interactive versions, all three pups were fully autonomous. In each run, the virtual camera followed the gray pup intelligently [Tomlinson 2000b], and tried to show off its interactions to greatest advantage.

The runs were randomized to account for possible order effects. The Non-interactive run was in the first position for 11 of the 32 subject, in the second position for 11 subjects, and in the third position for 10 subjects. The Button/AlphaWolf run was in the first position for 10 of the 32 subject, in the second position for 10 subjects, and in the third position for 12 subjects. The third, variable run was in the first position for 11 of the 32 subject, in the second position for 11 subjects, and in the third position for 10 subjects.

An error in the randomization methodology caused an uneven distribution within the variable run. The Microphone runs were distributed 5/0/3 (first run/second run/third run), The Button/Emotional runs were distributed 0/4/4, Button/Fixed were distributed 2/4/2, and Button/Random were distributed 4/3/1. I will address the extent to which this error confounds some of the results in section 5.2.3 below.

Each run lasted approximately four minutes. During the run, the system (unknownst to the subject), kept track of a variety of information about the subject’s interaction, e.g., how many button-presses, screen-clicks, etc. At the end of the run the system recorded these data, as well as the social relationships that had developed among the pups, to a file.

### **The Three Questionnaires**

Immediately after each run, the subject was asked to fill out a questionnaire. This questionnaire consisted of a field for age, a field for gender, 20 questions on a Likert scale [Likert 1932] from 1 to 7, and a final, open-ended question for “any additional thoughts.” The central questions addressed several main topics:

- Which wolf was more dominant?
- How much control did you have over the wolves?
- How similar was it to the real wolf video?
- How much did you enjoy the run?
- How much did you like the wolves?
- How much did you identify with the wolves?

The full questionnaire is attached at the end of this document (Appendix A).

Subjects were given as much time as they needed to fill out the questionnaire, and were instructed to knock on the door when they were finished.

### **Final Questionnaire**

After the three runs and three identical questionnaires, the subject was asked to fill out one final questionnaire (different from the first three), comparing the three runs. This questionnaire is attached at the end of this document (see Appendix). This questionnaire was different from the first three, in that it did direct comparisons of the three runs. There were 21 questions on a 7 point Likert scale, in which subjects did pairwise comparisons of the various runs. It required three questions to compare the three runs against each other for each topic (A vs. B, A vs. C, B vs. C). The seven areas of comparison were:

- Which run did you enjoy more?
- In which run did you feel you had more control over the behavior of the wolves?
- In which run did the behavior of the virtual wolves more closely resemble the behavior of real wolves?
- In which run did you feel more comfortable with the interface or lack thereof?
- In which run were the social dominance relationships clearer?
- In which run did you feel more immersed in the experience?
- In which run did you like the wolves more?

In addition to the 21 questions that addressed these seven topics, there was a final, open-ended question for “any additional thoughts.”

### **Debrief**

After the final questionnaire, the subject was given the promised ice cream gift certificate and read a short debriefing statement, informing them of the purpose of the study, the fact that one of their runs would be used anonymously in another subject’s experiment, and instructed them to notify the experimenter if they, now or later, suffered any physical or emotional ill effects as a result of their participation in the study. The subject was thanked, and directed to the elevators.

### **Process Data**

After the subject departed, the experimenter archived the consent form and input the data into a database. At this point, the experimenter began preparing for the next subject.

#### ***5.2.2.4 Subjects***

The 32 subjects were drawn primarily from the Cambridge and MIT communities. Subjects were recruited via emails to an assortment of mailing lists and posters placed around the Kendall Square area. The primary criteria for participation were that the subject had never interacted with the virtual wolves (either at SIGGRAPH or at the Media Lab, where they had been running for several months), and had never been exposed to the substance of the research behind the virtual wolf social behavior.

The subjects ranged in age from 17 to 55 (mean = 26.2, standard deviation = 7.8). Half (16) were female (min age = 18, max age = 55, average age = 26.1, standard deviation = 10.1), and half were male (min age = 17, max age = 37, average age = 26.4, standard deviation = 4.7). Twelve of the subjects were acquainted with the experimenter, and 20 were completely naïve.

There are some clear distinctions to be drawn among the subject base, irrespective of algorithms or interactions that each subject experienced. Both gender and age, in particular, had effects on subjects’ responses.

### **Gender Effects**

Female subjects gave slightly higher scores on subjective questions, regardless of the topic. Of the 32 subjects, 16 were male and 16 female. Questions 6-22 on each of the first three questionnaires could be seen as subjective (the first several questions pertaining to gender, age, and which pup was more dominant). The subjects therefore filled out a total of 1632 subjective questions (816 samples for each gender). On these questions, males gave an average score of 4.30, whereas females gave an average score of 4.64 ( $p < 0.0001$ ).

There are several possible explanations for this difference. First, it is possible that women are simply more willing to give a high score in response to value judgments.

A second possibility is that social relationship topics are more appealing to women. A way to tease apart this distinction would be to conduct another series of studies pertaining to a non-social topic, and yet keeping as much of the experimental method as possible the same.

A final possibility is that the experimenter was consistently male throughout the experiment, which may have introduced a bias among subjects. A way to clarify if this factor had any effect would be to conduct a follow-up study in which half of the subjects have a female experimenter.

Because of the randomization of the runs in the above study, it does not seem likely that the gender bias had a significant impact on any of the results.

### **Age Effects**

A second distinction among the population of subjects was an age bias. Subjects under 29 years of age gave slightly higher scores on subjective questions (6-22 on first questionnaires, average score 4.62) than those over 29 (average score 4.07,  $p = 0.0068$ ). These figures are based on a sample base of 23 subjects under 29 (1173 total subjective questions answered) vs. 9 users over 29 (459 total questions). The age of 29 (rather than the median value of 25) was chosen because it maximized the difference in average score between the two age groups.

Again, because subjects were randomized with regard to the runs they experienced and the order of the runs, it seems unlikely that this age bias had a significant effect on the results of the study.

#### ***5.2.2.5 The Virtual Wolves***

As described more fully in the Implementation chapter above, there were certain differences between the virtual wolves that people interacted with in the user tests versus those at SIGGRAPH. Primarily, various elements were removed from the SIGGRAPH installation to make the user-study wolves. There were no adults in the user tests, the pups did not age, and they did not play, howl or sleep. The interface was simpler, in that users could only direct their pups to growl or whine (rather than the full howl/growl/whine/bark control from the installation). These simplifications insured that the elements being tested were not confounded by confusion over too much complexity in too little time.

On the other hand, certain elements needed to be “played up”, since there was only one user interacting at a time in the user tests, rather than three at a time as there were in the installation. Instead of relying on other users to supply the autonomous behavior of the other pups, we instead turned to an autonomous behavior system that had been playing a secondary role in AlphaWolf. This behavior system caused pups to choose random other wolves to interact with, and to choose a style of interaction that matched their current emotional state.



In addition, so that relationships among the pups would be as clear as possible, we made the expression of the social relationships more extreme. While the internal representation remained the same, the animations tended to stay a bit closer to the extremes (where the hand-animated source material tended to be), thereby presenting a stronger but less subtle view into the wolves' social relationships.

### 5.2.3 Results

Thirty-two users each answered four questionnaires for a total of 2592 responses.<sup>4</sup> In addition, the system recorded the relationships for each litter of pups at the end of each run for a total of 576 relationships, each with a dominance value and a confidence value. Finally, the system tallied the total number of growls, whines, screen button clicks, wolf clicks and ground clicks that each user made each run (a total of 480 values). From these data, it has been possible to draw a variety of conclusions.

#### 5.2.3.1 Encoding, Decoding and Transmission

The first set of hypotheses is arguably the most important for validating the essential claim of this thesis – that the AlphaWolf mechanism encodes social relationships. In this section, I describe the structure of how I tested this premise, and present the results of the experiment.

Here is the exact statement of the hypotheses again:

**Hypothesis Group 1:** The AlphaWolf mechanism encodes social relationships created by a first person in a way that can be perceived by a second, naïve person. This hypothesis is, in fact, a combination of two sub-hypotheses: a) The AlphaWolf mechanism successfully encodes social relationships, and b) Naïve users can effectively perceive the relationships in our system.

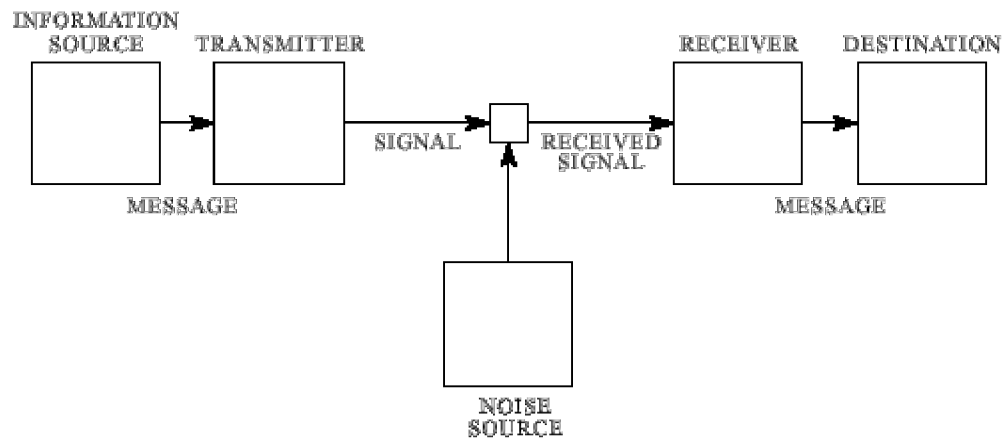


Figure 5-3: Schematic diagram of a general communication system (from [Shannon 1948]).

#### Shannon's Theory of Communication

In order to test the hypotheses above, I drew upon Shannon's Theory of Communication [Shannon 1948]. Shannon's theory subdivides the process of communication into 5 parts – an information source, a transmitter, a channel, a receiver, and a destination (see Figure 5-3). Each of these parts has a corresponding element in the AlphaWolf user tests. Each act of communication involved two subjects; for the sake of clarity, I'll call them Subject A and Subject B.

<sup>4</sup> One subject did miss one question, but it was part of a triad of questions from a final questionnaire. The other two questions in that triad suffice to determine a value for the missing answer.

B. Subject A's session with the wolves preceded that of Subject B; one of Subject A's Button runs created the relationships that were loaded into Subject B's Non-interactive run.

#### ***Information Source***

The first stage in the communication of social relationships is specifying the relationship to be transmitted. These relationships were specified in advance – the experimenter informed Subject A of the relationships that he or she should attempt to have the gray pup form with the white pup and the black pup. All subjects were instructed to be dominant to one pup and submissive to the other. Therefore, the experimenter is the information source, and the assigned relationships are the message that might or might not be transmitted to the destination.

#### ***Transmitter***

Subject A's interaction with the virtual wolf pack is the transmitter, taking the assigned relationships and converting them into the AlphaWolf dominance/confidence representation. This conversion encodes the relationships in a way that can be transmitted over the channel.

#### ***Channel***

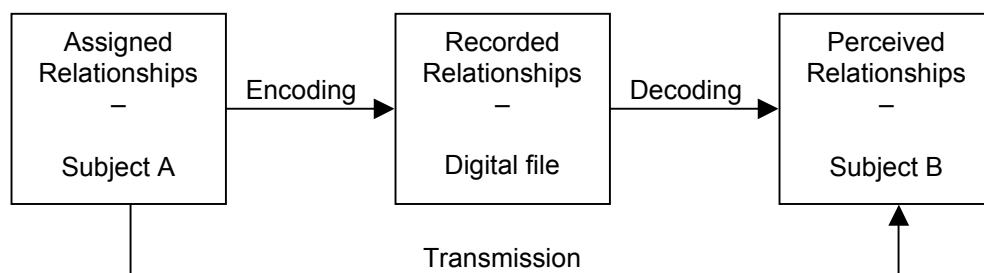
Whereas Shannon's theory deals primarily with a noisy channel (a telephone cable, for example), the channel in the experiment was a digital file that was loaded in without any loss of precision (barring experimenter error). This file stored the relationships formed by Subject A's interaction until they were loaded into Subject B's Non-interactive run.

#### ***Receiver***

Subject B's interaction serves as the receiver, decoding the social relationship representation into a human-readable form. (While the current representation is simple enough that it, too, is somewhat human-readable, a more elaborate version with three emotional axes, for example, would be completely opaque.) By viewing the social interactions among the virtual wolves in a Non-interactive run, Subject B attempts to recognize the relationships among the pups.

#### ***Destination***

The final destination of the message is the questionnaire that Subject B fills out. If Subject B reliably perceives the same relationships that were assigned to Subject A, and the only link between the two subjects is the AlphaWolf representation, then it is very likely that the representation successfully transmitted the relationships.



**Figure 5-4: Subject A is assigned the task of forming two social relationships. He or she encodes these relationships into a computational representation through an interactive run of the AlphaWolf system. These relationships are loaded in at the beginning of Subject B's Non-interactive run. If Subject B can perceive the relationships, they have been successfully decoded. If both encoding and decoding work, transmission occurs.**

The dissection of the process of communication in the above sections and in Figure 5-4 demonstrates that there might be difficulty at a variety of points in a single act of communication. In the experiment, Subject A might not understand the instructions. He or she might fail to encode the relationships because of a failure of the interface. There could be problems writing

the file at the end of Subject A's run, or loading it at the beginning of Subject B's. Subject B could fail to recognize the relationships. Subject B could get confused and fill out the form incorrectly. I constructed the experimental method in order to minimize the chance of each of the peripheral problems, so that I could focus on the two central issues at hand – whether Subject A can encode the social relationships, and whether Subject B can decode them. Only if both of these conditions are true can the AlphaWolf system be said to transmit social relationships.

### Encoding

Each of the 32 subjects viewed one Non-interactive run during their session. Thirty-one of the 32 subjects viewed Non-interactive runs based on other subjects' interactive runs. (The first subject had a Non-interactive run based on an interactive run from the pilot study, which had a slightly different setup. That subject will therefore not be included in certain portions of this analysis.)

Half of these runs were encoded by previous subjects interacting with runs that featured the AlphaWolf mechanism, while the other half were encoded by the Random algorithm or the Fixed algorithm. Of those encoded with the AlphaWolf mechanism, half again were done through the Emotional algorithm in which CSEMs are still formed, even though they are not allowed to influence a wolf's current emotional state. Each run included two assigned relationships, one between the gray pup and the white pup, and one between the gray pup and the black pup.

To determine whether or not encoding had occurred, I compared the assigned relationship to the recorded representation of the relationship. In order to convert the four-number representation of a dyadic relationship (dominance and confidence of wolf **A** towards wolf **B**, dominance and confidence of wolf **B** towards wolf **A**) into a form that could be compared readily to the assigned relationships ("dominant" or "submissive", from the point of view of the gray pup), I took an average of the dominance values, weighted by the confidence values. The relationship between two wolves (**A** and **B**) is determined by Equation 5-1:

$$\frac{((D_{AB} \times C_{AB}) + ((1 - D_{BA}) \times C_{BA}))}{(C_{AB} + C_{BA})}$$

**Equation 5-1**

where  $D_{AB}$  is **A**'s dominance value with respect to **B**,  $C_{AB}$  is **a**'s confidence in that dominance,  $D_{BA}$  is **B**'s dominance value with respect to **A**, and  $C_{BA}$  is **A**'s confidence in that dominance.

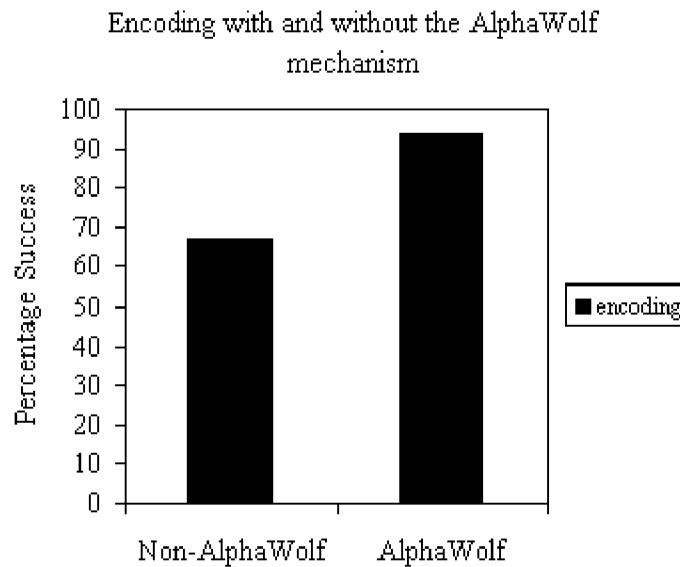
This formula yielded a single value from 0.0 to 1.0 that captured the essence of the relationship between the two wolves. If this value was less than 0.5, the relationship was "submissive" from the point of view of the gray pup. If it was greater than 0.5, the relationship was "dominant." If the relationship equaled exactly 0.5, it was called "middle." Because of the relative complexity of the system, it was exceedingly unlikely that a relationship would ever equal exactly 0.5. However, if two pups never met (as sometimes occurred between the black and the white pups, who might accidentally pass the four minute run without ever meeting), the value would end up at exactly 0.5 (a weighted average of the two starting values, 0.5 and 0.5). This condition did not happen in any of the assigned relationships involving the gray pup, though, and is therefore not relevant to these results.

Of the 32 total relationships encoded with the AlphaWolf mechanism (16 runs x 2 relationships per run), the internal representation matched the assigned relationship in 30 cases, and did not

match in 2 cases (93.8% success,  $p < 0.0001^5$ , see Figure 5-5). These figures demonstrate that the subjects and the system were quite successful at encoding social relationships.

Even in subjects' first run of virtual wolves, 15/16 were encoded successfully (93.8% success,  $p = 0.0021$ ). These figures demonstrate that there was not a significant learning curve in the experiment; people were just as successful at the beginning of the session as they were at the end.

In the runs without the AlphaWolf mechanism, in which relationships were either predetermined or random, 20 of the 30 runs were successfully encoded (66.7% success,  $p = 0.1261$ ). While this figure appears a bit higher than random (which predicts ~50%), the  $p$  value demonstrates that it is reasonable that it arose by chance. Nevertheless, it is well below the level of encoding shown by the AlphaWolf mechanism.



**Figure 5-5: The AlphaWolf mechanism offers significantly improved encoding over random chance.**

One 29-year-old male subject wrote a comment that suggests a possible reason the encoding success rate wasn't even higher. He reported a "[s]trong urge to test relationships – growl at black, for instance." The fact that people like to "poke at" a system is an important consideration in the design of interactive systems. Because the AlphaWolf system allows the computational entities to switch their relationships in response to adequate user input (e.g., it took about three clear interactions, in which both pups took on opposite roles, to reverse a relationship) the urge for people to vary their behavior could have had a significant impact on the results. Nevertheless, it appears that encoding was largely successful despite the possibility of intentional disobedience.

### Decoding

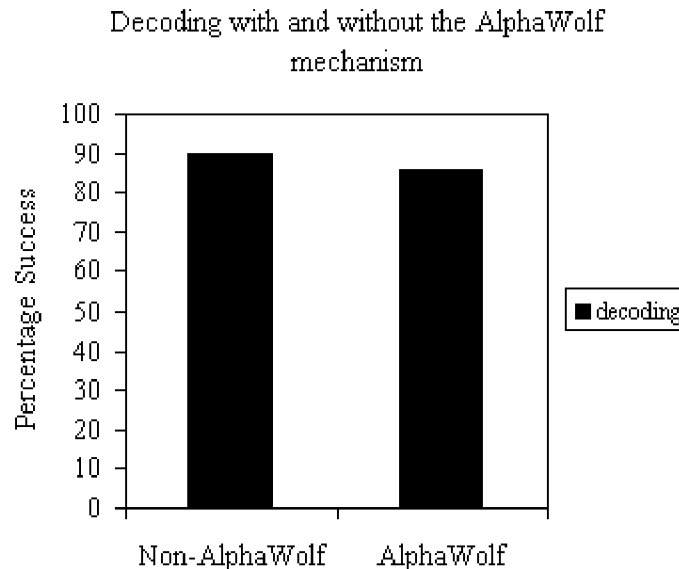
There were 31 subjects whose Non-interactive runs featured relationships created during another subject's run. Each of these Non-interactive runs included two relationships involving the gray pup (gray-black and gray-white). In the questionnaire at the end of that run, subjects were asked to specify which of each pair of pups was more dominant.

The way in which I determined a match in decoding was to simplify both the recorded representation and the subjects answers to the questionnaire. Using Equation 5-1 above, the

<sup>5</sup> Throughout this analysis, the Mann-Whitney U test was used to calculate the  $p$  values.

representation was reduced from four numbers to one, which in turn was converted to a simple “dominant” or “submissive” by the method described above. Subjects’ responses to the questionnaires were also simplified, with an answer of 1, 2 or 3 counting as “submissive”, an answer of 5, 6 or 7 counting as “dominant”, and an answer of 4 counting as “couldn’t tell.”

Of the 62 relationships (31 runs x 2 relationships per run), subjects were able to determine which pup was dominant in 53 cases, couldn’t tell in 2 cases<sup>6</sup>, and got it wrong in 7 cases (88.3% success,  $p < 0.0001$ , see Figure 5-6). These figures, well above chance levels, show that people were able to watch a run of virtual wolves and comprehend the relationships that they had seen there. Subjects were successful at decoding social relationships.



**Figure 5-6: Decoding is successful regardless of encoding mechanism.**

Even in their first run, subjects were successful at “reading” the relationships among the pups; of 22 first run relationships, subjects decoded 18 correctly, couldn’t tell for 1, and chose incorrectly in 3 (85.7% success,  $p = 0.0048$ ). These figures demonstrate that fully naïve subjects were able to determine relationships from only their pre-existing knowledge and their viewing of real wolves; they did not just learn what our system “meant” by social relationships over the course of their interactive runs.

There was no significant difference between decoding success when the relationships had been encoded by the AlphaWolf mechanism versus when they had been encoded by a random or fixed mechanism. Of the 32 decoded from AlphaWolf-created relationships, 26 were correct, 2 couldn’t tell, and 4 were incorrect (86.7% success,  $p < 0.0001$ ). Of the 30 decoded from non-AlphaWolf origins, 27 were correct and 3 were incorrect (90.0% success,  $p < 0.0001$ ).

#### ***The Relationships between Black and White***

Although the black and white pups were forming relationships with each other as well as with the gray pup, these pups were not the camera system’s “lead actor”, and therefore it was not very common for the camera to capture a clear dominance interaction between the two. As one 30-year-old male subject put it: “It was hard to determine the relationship between the two other pups (in this case between white & black) as I rarely saw them interacting with each other (only

<sup>6</sup> For the purposes of calculating averages and p values, runs in which subjects couldn’t tell the relationship were discarded.

at the end).” Added to which, only the two relationships involving the gray wolf were assigned to the encoding subjects. Therefore, the relationships between black and white were not included in the above results for decoding.

Nevertheless, it appears that subjects did just as well decoding the relationships between these two pups as they did in the relationships involving the gray pup. Of the 32 black/white relationships, subjects successfully decoded 26, subjects couldn’t tell in 2 cases, got it wrong in 2, and in 2 cases the black and white pups never met<sup>7</sup> (and therefore had no relationship) (92.9% success,  $p < 0.0001$ ). Even in their first run, subjects were successful; of the 11 subjects whose Non-interactive run was in the first position, there were 10 correct decodings, 1 run where the pups did not form relationships, and no mistakes. (100% success,  $p = 0.0079$ ) It appears that it doesn’t take much screen-time for people to recognize a relationship!

In a triadic dominance system such as the one among these pups, there could be another reason why subjects performed so well with so little direct information. If decoding subjects have a strong view of gray’s relationship with black and a strong view of gray’s relationship with white, and if the encoding subjects were always assigned to play gray, to dominate one pup, and to submit to the other, then, by the transitive property of social relationships, subjects might be able to predict the relationship between black and white. Nevertheless, packs encoded with the random algorithm would arguably detract from this phenomenon, and pups encoded during a Fixed run would have strong relationships, but would not be guaranteed to have gray in the middle rank. Despite these facts, it is relevant that certain “whole-world” effects may have contributed to people achieving some indirect understanding of relationships to which they were not well-exposed directly.

### Transmission

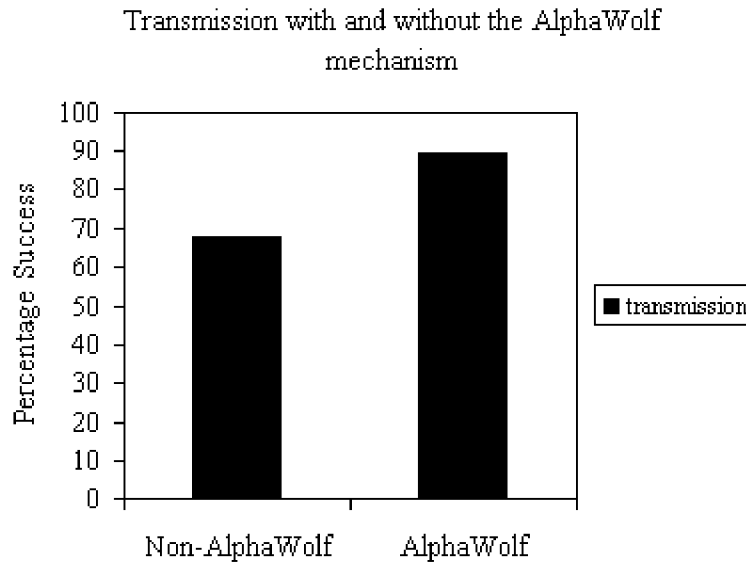
Considering the entire process (both encoding and decoding) as a single act of communication, there were 16 subjects whose Non-interactive runs were based on another subject’s interactive run that featured the AlphaWolf algorithm. Each of these subjects’ Non-interactive runs had two relationships, for a total of 32 relationships. Of these, 26 were successfully transmitted (assigned relationship matched recorded representation and answer given on questionnaire), 1 was a double error resulting in successful communication of the message (assigned relationship matched answer given on questionnaire, but recorded relationship was opposite<sup>8</sup>), 2 couldn’t tell, and 3 were incorrect (89.7% success,  $p < 0.0001$ ). These figures verify that the AlphaWolf algorithm and the entire AlphaWolf system successfully captures, transmits and displays social relationships in a way that is readily apparent to humans.

---

<sup>7</sup> For the purposes of calculating averages and p values, runs in which pups never formed a relationship were discarded.

<sup>8</sup> Throughout this analysis, double errors were thrown out in determining percentages and p values. Nevertheless, it is an interesting possibility that “whole-world” effects somehow contribute to the prevalence of double errors – that some aspect of the inter-relation among the three pups makes people able to divine the correct relationship despite a flawed representation.

In this particular double error, the recorded relationships suggested that both pups had decided that they were submissive to the other. This unusual case most likely created behavior that was difficult to understand for the perceiving subject. The subject gave a response of 5, suggesting a slight leaning towards gray being dominant. In fact, gray *was* slightly more dominant, feeling 0.12 dominance towards white, while white felt 0.08 dominant towards gray. However, gray had a higher confidence in his submissive state (0.87 vs. white’s 0.65), so the weighted average showed a slight lean towards gray being more submissive (0.46).



**Figure 5-7: Transmission is significantly better with the AlphaWolf mechanism than random chance.**

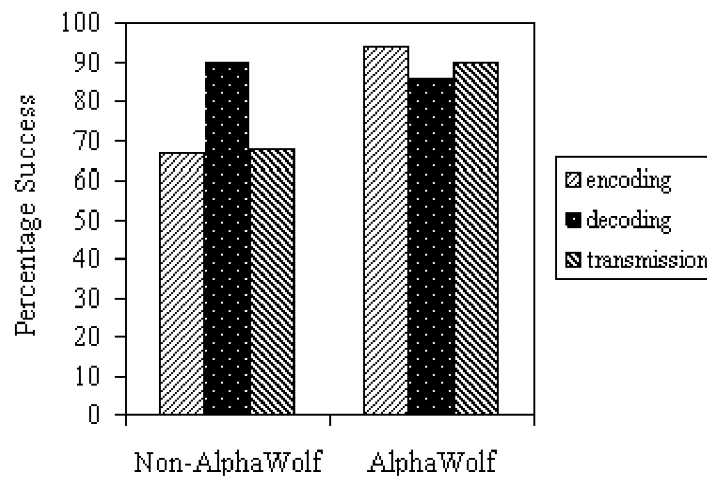
The subjects who viewed Non-interactive runs for which the relationships had not been encoded by the AlphaWolf mechanism performed significantly less well. Of the 30 relationships, transmission occurred in 19 cases, there were 2 double errors, and 9 incorrect responses (67.9% success,  $p = 0.0574$ ). These figures are surprisingly high, but well below the level attained by the AlphaWolf mechanism.

### Discussion

To summarize the above study: 32 human subjects interacted with packs of virtual wolves under controlled conditions. Each person played the role of a gray wolf pup and attempting to direct that pup to form certain dominance relationships with its two siblings. At the end of the run (approx. 4 minutes long), the system recorded all of the pups' internal representation of the relationships. A second experimental subject later viewed a Non-interactive pack of autonomous wolves whose relationships were specified by the recorded relationships from the previous subject's interactive run. If the second subject's perception of the relationships matched the relationships that had been assigned to the first subject, then the AlphaWolf system succeeded in encoding, transmitting and decoding those relationships. This method of analysis is derived from Shannon's Theory of Communication (1948).

The results from that study demonstrate that the AlphaWolf mechanism successfully represents social relationships (see Figure 5-8). Using the AlphaWolf mechanism, the first subjects successfully encoded 93.8% of the relationships that they had been assigned to create ( $p < 0.0001$ ). The second subjects decoded 88.3% of the relationships that they viewed ( $p < 0.0001$ ). These two factors resulted in a successful transmission rate of 89.7% ( $p < 0.0001$ ). Each of these figures demonstrates significantly better performance than the 50% success rate predicted by chance. The significantly lower levels of encoding and transmission among non-AlphaWolf runs confirms that these effects are a direct result of the AlphaWolf algorithm.

### Encoding, Decoding and Transmission with and without the AlphaWolf mechanism



**Figure 5-8: The AlphaWolf mechanism offers significantly better encoding, and therefore transmission, than chance (50%) or than a system without the AlphaWolf mechanism. Subjects were successful at decoding relationships regardless of how they were encoded.**

#### *Black and White*

To offer an argument in opposition to the validity of this experiment's results, one could say that, since subjects were only given limited controls, and were directed in how to dominate or submit (press one of the two buttons), they were not really encoding relationships, but simply obeying instructions. Imagine this hypothetical experiment: rather than virtual wolves, subjects were presented with a gray screen, and were asked to press one button to make it blacker and another to make it whiter. At the end of the interaction, the representation stored the brightness of the screen. A second subject was presented with a screen at that brightness, and asked to fill out a questionnaire specifying if the screen was black or white. It would be unsurprising if all of the subjects got it right.

What this thought experiment points out, though, is that social relationships are almost as clear to people as black and white. Dominance and submission are the endpoints on an axis about which people (and presumably wolves) naturally think. The AlphaWolf mechanism captures social relationships from dominant to submissive as well as a single number can represent brightness from black to white. In addition, it confirms that the supporting animations, behavior systems, camera systems, interfaces, perception systems, and installation design components do not significantly hinder the central mechanism's ability to function.

Nevertheless, a revised experimental protocol might involve giving the encoding subject several unlabeled keys, and allowing them to work out which keys make the wolf behave in ways that appear to them to be dominant.

#### **5.2.3.2 Effects of Social Relationship Algorithms**

The second main area of experimentation in the user studies involved comparing the effect of various algorithms for social relationship formation on the subjective experience of an interactor. Here is the statement of the hypotheses again:



**Hypothesis Group 2:** A run with the AlphaWolf algorithm will outperform a run with the Emotional algorithm, which will in turn outperform a run with the Fixed algorithm, which will in turn outperform a run with the Random algorithm (i.e. AlphaWolf>Emotional>Fixed>Random), in each of the following areas: similarity to real wolves, clarity of social relationships, user enjoyment, user control, user comfort with the interface, user liking of the wolves, user identification with the wolves, and user immersion.

### Experimental Method

In order to test these hypotheses, 24 subjects interacted with runs of virtual wolves that were identical in all respects except for the underlying algorithm (and the order in which they were presented, which I will address later.) All of these subjects interacted with a Button/AlphaWolf run. Eight of the subjects also had a Button/Emotional run. Eight others had Button/Fixed. Eight had Button/Random. (The remaining eight of the 32 total subjects had Microphone/AlphaWolf, which will be addressed in the section 5.2.3.3.)

On each of the runs, subjects were read an identical set of directions. In addition, each subject was assigned to create the same social relationships on the two runs<sup>9</sup>. Half of the subjects were instructed to dominate the black pup and submit to the white pup; for the other half, the relationships were reverse. Within each subject's session, the assigned relationships were held constant, so that they would not bias the results.

As described above, each subject filled out a questionnaire after each run answering a variety of questions about that run. At the end of the session, the subject filled out an additional questionnaire comparing the three runs. Elements of the data below are derived from both interim questionnaires and from the final questionnaire. Runs could be compared both by means of analyzing the absolute scores given to each run in the interim questionnaires and by means of the direct comparisons ranked in the final questionnaire.

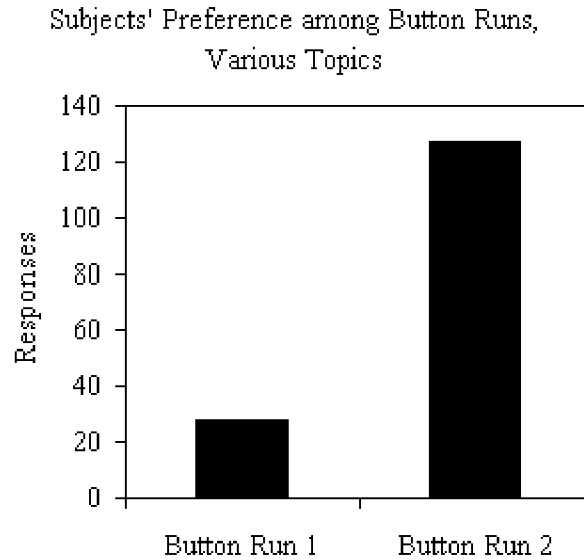
### Results

The results in this section appear to have been swamped by the effect of the order of the runs. Regardless of algorithm, subjects showed a strong preference for the second of two Button runs. Each of the 24 subjects who had two Button runs answered 7 questions comparing those runs. Of these 168 responses, fully 81.9% of those which expressed a preference chose the second run (127 vs. 28, with 13 ties,  $p < 0.0001$ , see Figure 5-9)<sup>10</sup>. This preference appeared to hold regardless of the algorithm being compared (41/9/6 AlphaWolf vs. Emotional (secondRun/FirstRun/Ties), 40/14/2 AlphaWolf vs. Fixed, 46/5/5 AlphaWolf vs. Random). Also, the preference held regardless of topic ("enjoyment" 19/5/0, "control" 18/5/1, "like wolves", 19/1/4, "comfortable interface" 19/3/2, "clearest relationships" 16/5/3, "immersing" 20/4/0, "likeable wolves" 16/5/3). Therefore, there was sparse evidence for the effect of the various algorithms on any of the factors.

---

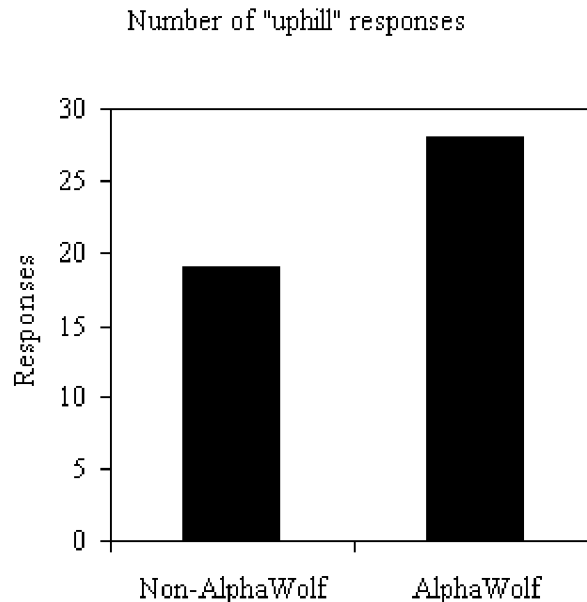
<sup>9</sup> The decision to have each subject form identical relationships on each interactive run was the main conclusion reached during the pilot study. The effect of a subject taking on different roles (i.e., dominate both, submit to both, or submit to one and dominate the other) appeared to be swamping the effects that we were trying to test.

<sup>10</sup> P values in this section determined with binomial test.



**Figure 5-9: Subjects preferred the second Button Run, regardless of algorithm.**

Nevertheless, of the 28 answers that were “uphill” (i.e., where subjects chose the first run, against the strong universal favoring of the second run), 19 were for AlphaWolf (67.9%,  $p = 0.0436$ ). Of these uphill votes, in the category of “most control” AlphaWolf won 4/5, in “clearest relationships” it won 4/5, and in “most likeable wolves” it won 4/5. These data offer a reasonable likelihood of AlphaWolf having the edge over the three other algorithms on a variety of measures.



**Figure 5-10: AlphaWolf was favored with more responses than the other algorithms were when a first Button run was chosen over a second Button run.**

### Discussion

Despite the lack of clear data separating the four individual algorithms, the fact that people exhibited such a strong preference for the second run better is a good sign for the installation as a

whole. It suggests, at least, that subjects hadn't gotten tired of it after the first run. This theory is backed up by several comments that subjects wrote in at the end of questionnaires.

- "I think I was just learning on Button Run I." (26-year-old female)
- "A practice run with the navigation would have been helpful to get acquainted with the camera pan/zoom and interaction with the horizon. (I don't play video games.) ... I was too busy experiencing it for the first time on Button I." (37-year-old male)
- "It was very cool! I enjoyed it! I didn't remember to pay attention to the wolves' initial behavior toward each other. I will next time." (20-year-old female)

So if people were "just getting the hang of it" during the first run, it is unsurprising that they would like the second run better.

Despite the predominance of subjects' taking a little while to get up to speed, a few subjects thought that the interaction was too simplistic. A 21-year-old male subject wrote, after his first interactive run: "Interesting at first... behaviors got rather predictable."

A possible change to the experimental method that could help minimize the order effects would have been to allow subjects to "play around" with the wolves until they feel comfortable with the interface, before starting the main experimental runs. This change would accommodate both the confirmed computer-gamers, who picked up the interaction paradigm with ease, and those subjects who needed a little more time to master the interaction.

#### ***A Preference for Emotion?***

Due to the randomization error mentioned above, there is another possible explanation for some part of the apparent order effects. The Emotional algorithm ended up being stacked in the later runs (none of the first runs were Button/Emotional, 4 of the second runs were, and 4 of the third runs were), with 6 of the 8 runs coming after the subject's Button/AlphaWolf run. Subjects showing a marked preference for the Emotional runs; of the 56 questions comparisons between AlphaWolf runs and Emotional runs (8 subjects x 7 categories), 41 chose the Emotional run, 9 chose the AlphaWolf run, and 6 marked it a tie (82% for Emotional,  $p < 0.0001$ ).

It is possible that subjects simply like the Emotional algorithm better, and this preference makes it appear that they favored later runs. It wouldn't be terribly surprising if this were the case, since the Emotional algorithm arguably does a better job of exploring the entire emotional repertoire of the pups, causing them to spend much of their time at the hand-animated end-points in their expressive dynamic range. This strong, clear, well-animated emotional expressiveness may well account for some of the favoring of later runs, since more of those runs featured the Emotional algorithm.

Nevertheless, even discounting subjects who had Emotional runs and looking at only subjects with Button/Random or Button/Fixed runs, there was still a strong order effect. Among 16 users responding on the 7 topics (112 total responses), 86 exhibited a preference for the second of the two button runs, 19 chose the first run, and 7 scored a tie (81.0%,  $p < 0.0001$ ).

Therefore, it appears likely that the order effects were much stronger than the perceivable differences between the algorithms, and the only appreciable result comes from those who voted "uphill", choosing the first Button run over the second. Nevertheless, further study would be necessary to clarify the situation.

### 5.2.3.3 *Interaction Paradigms*

The third area of experimentation in the user studies compared the effect of various interaction paradigms on the subjective experience of the interactors. The three interaction paradigms that were tested were a) pressing keys on a keyboard and moving a mouse to direct the gray wolf (Button), growling or whining into a microphone and moving a mouse to direct the wolf (Microphone), and a Non-interactive version. Here is the statement of the hypotheses that were mentioned earlier:

**Hypothesis Group 3:** A run with a microphone interface will outperform a run with a button interface, which will in turn outperform a non-interactive run (i.e., Microphone>Button>Non-interactive), in each of the following areas: similarity to real wolves, clarity of social relationships, user enjoyment, user control, user comfort with the interface, user liking of the wolves, user identification with the wolves, and user immersion.

### **Experimental Method**

The experimental method for this section was very similar to the experimental method used for Hypothesis Group 2. Each of the eight subjects in this dataset had one Non-interactive run, one Button/AlphaWolf run, and one Microphone/AlphaWolf run, in random order.

There was a slight difference between the Microphone run and the Button run because they required different instructions. The instructions for both are included in Appendix 0. Aside from this necessary variation, everything about the two runs was held constant.

### **Results**

The results for this set of hypotheses were much clearer than those for Hypothesis Group 2. Since all 32 subjects had one Non-interactive run and two interactive runs, there are ample data for comparing interactive and non-interactive versions of the virtual wolves. With only 8 subjects having Microphone runs, the data are a bit sparser for distinguishing between Button and Microphone interaction paradigms. Nevertheless, there are some interesting conclusions to be drawn from the data.

The order of presentation of the different runs did not appear to have a significant effect on subjects' answers with regard to the interaction paradigms. In fact, whereas the previous data show a pronounced preference for later runs, the data for the eight users who had a Microphone run show a slight preference for earlier runs. In the 56 final questionnaire responses comparing the Button run with the Microphone run (8 users x 7 categories), 34 favored the first run, 19 chose the second run, and 3 voted a tie (64.2% in favor of first run,  $p = 0.0267$ ) While the probability of a strong subjective preference for later runs of the wolves (from the previous section comparing two Button runs) makes it hard to establish the exact magnitude of effects in this section, the fact that subjects showed a slight preference the other way suggests that they were scoring runs based on the interaction paradigm, rather than simply on the order of presentation; therefore the qualitative component of the results should be valid. These data suggest that the effect of the interaction paradigm is greater than the effect of order, which is, in turn, greater than the effect of social relationship algorithm on the subjective experience.

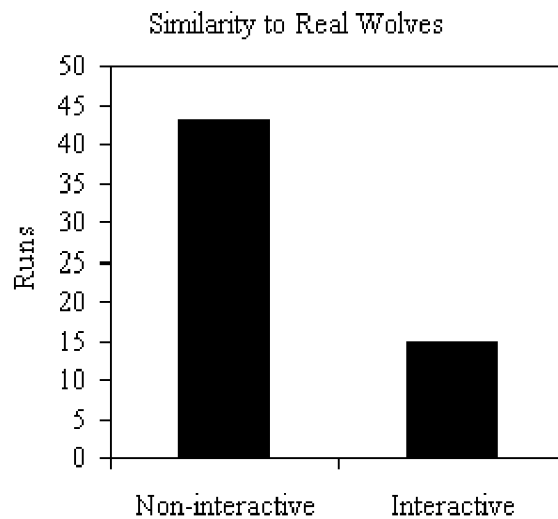
In the next eight sections, I address each of the areas that were being explored for effects caused by variation in interaction paradigm. In most of these sections, I first compare Non-interactive runs to the other two runs, using results from all 32 subjects. Each subject answered three questions in the final questionnaire pertaining to each of these seven topics. Each question did a

pair-wise comparison of two runs. These questions form the basis of the results below. I also support several of the conclusions with data from the questionnaires about each individual run.

### ***Similarity to Real Wolves***

The first area is which kind of interaction makes people feel that the virtual wolves are most similar to the real wolves that they watched in the National Geographic Video at the beginning of their session. The question that subjects were asked was: “In which of the runs did the behavior of the virtual wolves more closely resemble the behavior of real wolves?”

Subjects demonstrated clearly that non-interactive wolves are more like the real thing than runs in which they were asked to interact. Of the 64 questions comparing a Non-interactive run to another run on this topic, 43 responses chose the Non-interactive run, 15 chose the other run, and 6 were ties<sup>11</sup> (74.1%,  $p = 0.0002$ , see Figure 5-11). As a 21-year-old female subject put it: “It seems to me, that when I was watching the video, that seemed more like ‘real’ wolves to me, insofar as I didn’t have control over them. I also have no control really over a real wolf.”



**Figure 5-11: Subjects found the pups in Non-interactive runs to be more like real wolves than those in interactive runs.**

This conclusion is supported as well by the answers given to the questionnaires administered after each run. Across all Non-interactive runs (32) and all interactive runs (64), the Non-interactive runs received consistently higher scores than the interactive runs on “similarity of overall social behavior” (5.3 vs. 4.9,  $p = 0.2134$ ), “similarity of submissive behavior” (5.7 vs. 5.3,  $p = 0.1852$ ), and “similarity of dominance behavior” (5.5 vs. 4.8,  $p = 0.0095$ ).

The results comparing Button to Microphone interaction paradigms were inconclusive. Six of the 8 subjects with one of each kind of run chose Button as being more like real wolves (75%,  $p = 0.1445$ ). Nevertheless, in this same set of 8 subjects, two of the three areas addressed by the initial questionnaires came down on the side of the Microphone run: “similarity of overall social behavior” (Button 5.0 vs. Microphone 5.3,  $p = 0.8785$ ), “similarity of submissive behavior” (5.4 vs. 4.9,  $p = 0.5054$ ), and “similarity of dominance behavior” (4.9 vs. 5.0,  $p = 0.9591$ ). One 44-year-old female subject had the following to say on the topic: “I was very preoccupied w/ dual use of microphone + mouse (self-conscious; nervous) +, frankly, lost sight of virtual wolves behavior as compared with natural wolves + with each other, but the graphics were engaging.”

<sup>11</sup> Ties are discarded for the purposes of percentages and p-values.

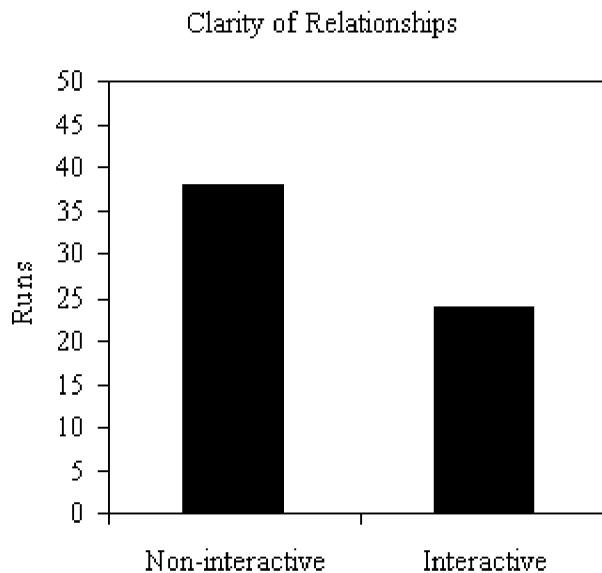
As I mentioned earlier, perhaps a warm-up run would have helped to clarify the effect of the interface.

To summarize, non-interactive wolves act more like real wolves than interactive wolves, but Button vs. Microphone did not have an appreciable effect.

### ***Clarity of Relationships***

The second area investigated was the clarity of the relationships between the virtual wolves. The question asked was: “In which of the runs were the social dominance relationships clearer to you?”

Subjects agreed that Non-interactive runs offered a clearer view of the wolves relationships than interactive runs. Of the 64 questions comparing a Non-interactive run to an interactive run in this category, 38 chose the Non-interactive run, 24 chose the other run, and 2 were ties (61.3%,  $p = 0.0490$ , see Figure 5-12). As one 30-year-old male subject mentioned, “I was so busy dominating and submitting that I didn’t think to observe who was dominant btw. white & black until right before the end.” Again, giving people a chance to get used to the interaction paradigm before the trial runs might have helped to clarify these results.



**Figure 5-12: Relationships are clearer in Non-interactive runs.**

This result was supported (although not significantly) by responses given in the first three questionnaires to the question “How clear were the social dominance relationships between the virtual wolves?” Non-interactive runs scored an average of 5.28, while interactive runs scored a 5.18 ( $p = 0.6173$ ). While this difference certainly wouldn’t be convincing on its own, it does support the stronger expression of the same tendency in the final questionnaires.

With regard to Button versus Microphone interaction paradigms, six of the eight subjects preferred the Button run for clarity of relationships (75%,  $p = 0.1445$ ). This result is contradicted slightly by the scores given to the corresponding question in the initial questionnaires – among the eight subjects with Microphone runs, Button had an average score of 4.75, and Microphone runs scored 4.875 ( $p = 0.8785$ ). Again these results are not significant.

To summarize, non-interactive runs have clearer relationships than interactive runs, but it was not possible to distinguish the effect of the Button vs. the Microphone interaction paradigm. This clarity could result from subjects' ability to concentrate on the relationships when they are not interacting, or also from the lack of "competition" between the gradually forming personality of the semi-autonomous pup and the actions directed by the participant.

### ***User Enjoyment***

The third area of comparison is the enjoyment that subjects felt in the interaction. The question asked was: "Which of the runs did you enjoy more?"

In their responses to the final questionnaire, subjects did not appear to enjoy the Non-interactive runs as much as they enjoyed the interactive runs, though without statistical significance. Of the 64 responses, 35 chose the interactive run, 24 chose the Non-interactive run, and 5 scored a tie (59.3%,  $p = 0.0963$ )

This result is contradicted slightly (albeit insignificantly) by subjects' responses to the initial questionnaires. Subjects gave an average score of 5.2 to Non-interactive runs and a score of 5.1 to interactive runs ( $p = 0.9554$ ).

Subjects did not show a strong preference one way or the other between Button and Microphone, either, in terms of enjoyment. Of the eight subjects with one of each kind of run, three chose the microphone, four chose Button, and one gave them a tie (57% for Button,  $p = 0.5000$ ). The lack of preference is mirrored by the responses given on the initial questionnaires – among these 8 subjects, Button received an average score of 5.0, and Microphone runs received a 4.8 ( $p = .4418$ ).

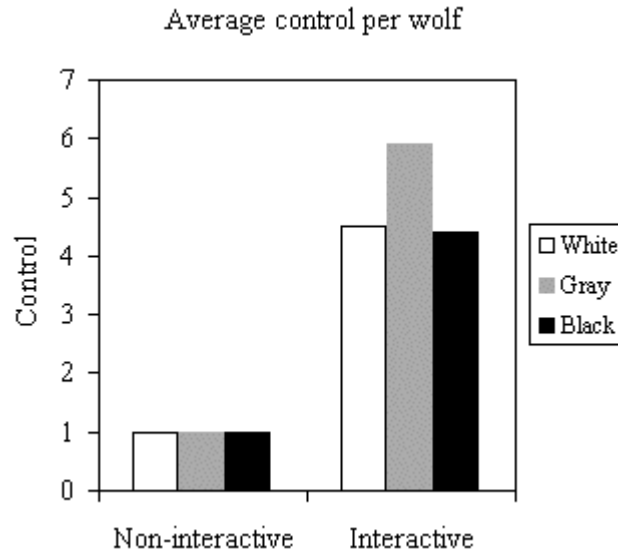
To summarize, there was no clear distinction among the three interaction paradigms with regard to user enjoyment.

### ***Control***

The fourth area of comparison is the amount of control that subjects feel they have over the run. The question asked was: "In which of the runs did you feel you had more control over the behavior of the wolves?"

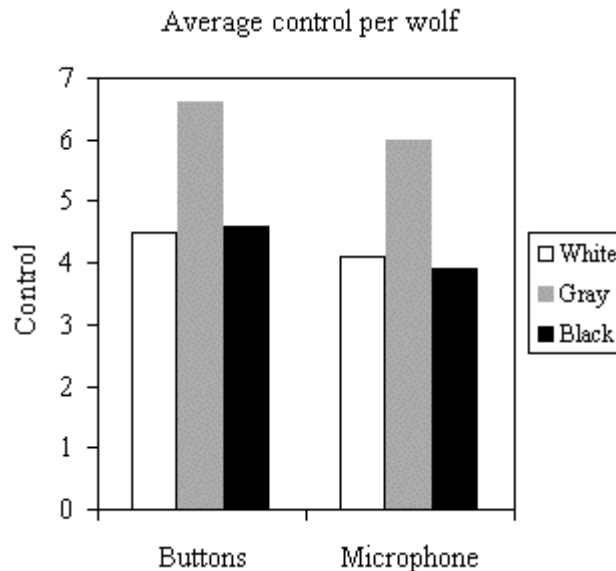
Unsurprisingly, Non-interactive runs lost this one in a landslide. Of the 64 questions, Non-interactive won only 2, and lost the other 62 (96.9%,  $p < 0.0001$ ). The comforting part of this clearly polarized result is that it confirms that subjects were, for the most part, understanding the directions and answering honestly.

This result is confirmed strongly by the responses to the initial questionnaires. Across all interactive runs (64) and Non-interactive (32) runs, subjects felt more control over the wolves' relationships (4.8 vs. 1.0,  $p < 0.0001$ ), 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,  $p < 0.0001$ ). The fact that people felt more control over the gray pup than over black or white is not surprising, since that was the pup over whom they had direct control. Nevertheless, it confirms that everything was essentially "working" in the runs.



**Figure 5-13: Subjects felt more control in the interactive runs.**

Between the Button and Microphone runs, people appeared to feel more control with the Button. As one 24-year-old male subject wrote: “I found the buttons much easier to use than the microphone (I’m not too good at making animal sounds!)” Of the 8 people with both kinds of run, 6 of them chose the Button run (75%,  $p = 0.1445$ ) with regard to this topic. This preference was mirrored by the subjects’ responses to the initial questionnaires as well. Button outsourced Microphone 5.3 to 4.4 for control over the wolves’ relationships ( $p = 0.2345$ ), 4.5 to 4.1 for control of white’s behavior ( $p = 0.5737$ ), 6.6 to 6.0 for control of gray’s behavior ( $p = 0.1304$ ), and 4.6 to 3.9 for control of black’s behavior ( $p = 0.3828$ ). These data are presented in Figure 5-14.



**Figure 5-14: Subjects felt slightly more control with buttons than with the microphone.**

To summarize, the Button interaction paradigm appeared to provide a bit more control than the Microphone interaction paradigm, in particular over the behavior of the gray pup. It is unclear whether the preference for buttons in this section is the result of the greater ease of use of the



buttons or some other factor. Nevertheless, both buttons and microphone were significantly superior to the Non-interactive runs in terms of control.

### ***Comfort***

The fifth area of comparison is how comfortable people found the various interfaces to be. The question asked was: “In which of the runs did you feel more comfortable with the interface or lack thereof?”

This area was won, weakly, by the interactive runs. Of 64 responses, the interactive run won 34, Non-interactive won 26, and they tied in 4 cases (56.7%,  $p = 0.1831$ ). There was no question on the initial questionnaires to support or deny this slight preference.

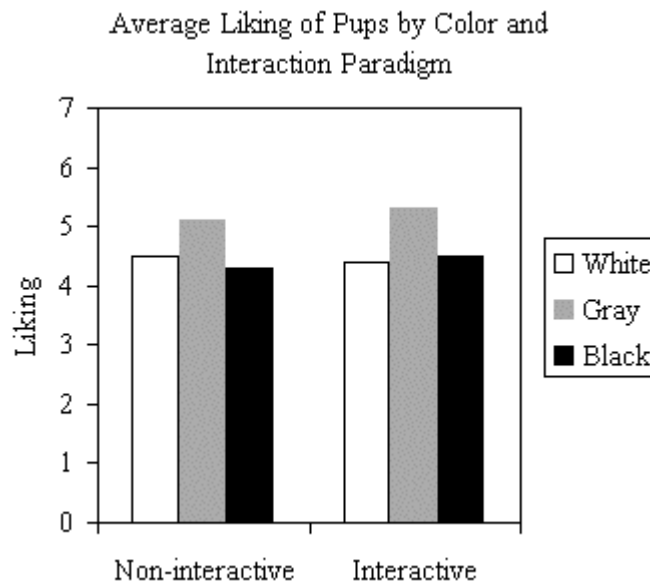
Between Button and Microphone, five of the eight subjects with both kinds of run chose Button as the more comfortable interface (63%,  $p = 0.3633$ ).

To summarize, although Non-interactive scored slightly higher than interactive runs, and Button scored slightly higher than Microphone, neither result was statistically significant.

### ***Liking of Wolves***

The sixth area of comparison is how much people liked the virtual wolves themselves. The question asked was: “In which of the runs did you like the wolves more?”

Subjects appeared to like the wolves better in Non-interactive runs than in interactive runs. Of the 64 pair-wise comparisons, Non-interactive won 35, lost 24, and tied 5 (59.3%,  $p = 0.0963$ ).



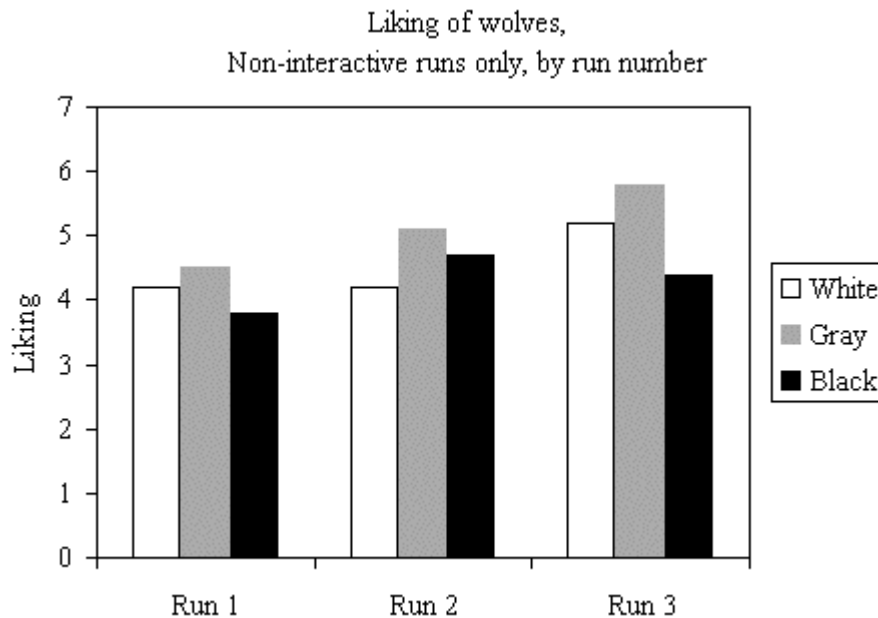
**Figure 5-15: Cinematography makes more of a difference than interactivity in how much people like pups.**

In the initial questionnaires, however, the results were inconclusive at distinguishing interactive from non-interactive runs. In terms of how much they liked the whole pack, subjects gave the interactive packs a score of 5.0 versus 5.2 for Non-interactive packs ( $p = 0.7348$ ). Each pup was just about evenly matched across interaction paradigms, too, with white receiving scores of 4.4 on interactive runs and 4.5 on non interactive runs ( $p = 0.6321$ ), gray receiving scores of 5.3 and 5.1

respectively ( $p = 0.5380$ ), and black receiving 4.5 and 4.3 respectively ( $p = 0.6311$ ). Figure 5-15 shows these data.

The most apparent result here is that subjects liked the gray pup more than the black or white pups regardless of interaction paradigm (gray average score across 96 runs = 5.2, black and white average score = 4.4,  $p = 0.0001$ ). This effect appears to result to a large extent from the camera's focus on the gray pup, and to a lesser extent from the control that people have over the gray pup, since the effect is slightly less pronounced in the Non-interactive runs (5.1 vs. 4.4,  $p = 0.0644$ ) than it is in the interactive runs (5.3 vs. 4.5,  $p = 0.0005$ ), even though the camera algorithm was identical in both.

A possible explanation for the apparent significance of cinematography could result from the cross-over among the various runs, since gray was the lead actor or the directed character in every run. For example, the fact that a subject played the role of gray could explain part of why he or she continued to like gray in a later, non-interactive run. This possibility is undermined by the fact that there was a clear preference for the gray pup even in Non-interactive runs in the first position, when subjects had no knowledge of which pup they might eventually play (see Figure 5-16). In the first run, gray scored a 4.5, versus white's 4.2 and black's 3.8. Nevertheless, later Non-interactive runs did show a stronger separation; in the first run gray outscored the other pups by an average of 0.5, in the second run the separation was 0.6, and in the third run the split was 1.0. (For the sake of comparison, the interactive runs showed a disparity of 0.6 in the first run.) Since the preference for gray did get greater in later runs, the cross-over effect may account for some of the effect.



**Figure 5-16: People tend to like the wolves more in later Non-interactive runs than earlier ones.**

Between the Button runs and the Microphone runs, subjects exhibited a slightly greater liking of the wolves in Button packs. Among the 8 subjects with both a Microphone run, the average liking of the entire pack was 5.3 in Button runs versus 4.4 in Microphone runs ( $p = 0.1304$ ). Their liking of the white pup was 5.0 and 4.5, respectively ( $p = 0.1949$ ). Their liking of the gray pup was 5.1 and 4.6, respectively ( $p = 0.4418$ ). Their liking of the black pup was 4.3 and 4.4, respectively ( $p = 0.9591$ ). While the sample size was insufficient to achieve statistical significance, it appears that people may have liked the wolves better when interacting with Button

than when interacting with the Microphone (4.8 average score for the three pups vs. 4.5,  $p = 0.2450$ ).

A possible alternative explanation of people's preferential liking of the gray wolf lies in the fact that most of the gray wolves were playing the middle rank of the hierarchy. Gray was in the middle on all of the interactive runs (or at least, the user had been assigned to give it the middle role), and in most of the Non-interactive runs as well, since at least half of these runs were based on successfully encoded runs, which put the gray pup in the middle. It is possible that people tended to strongly dislike either the very dominant or the very submissive pup, and the combination of an even split between these two preferences left gray winning, for lack of any strong dislike.

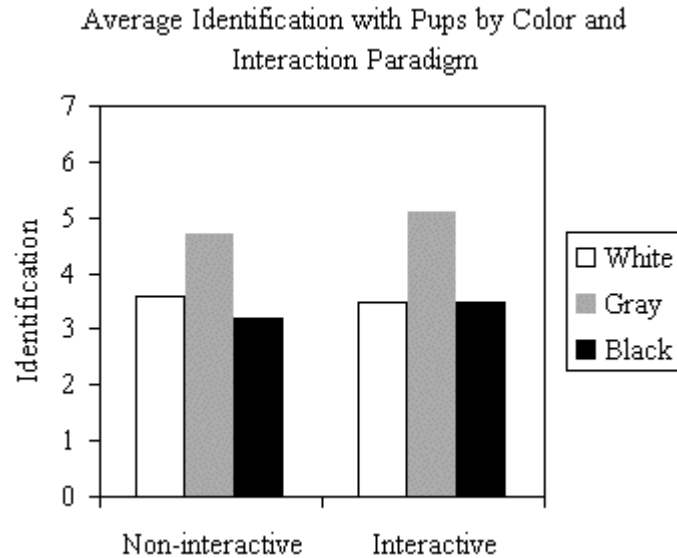
It is also possible that this effect was simply caused by people "liking the look" of the gray pup more; however, it was for this reason that the gray pup, in the center of the brightness spectrum, was chosen as the lead actor; I feared that the black and white pups might have too much built-in meaning simply on account of their coloration. Nevertheless, it seems far more likely that cinematography lies at the root of the issue.

To summarize, subjects appear to like the wolves about the same amount in interactive and non-interactive runs, with a slightly greater preference in Button runs over Microphone runs. They exhibited a clear preference for the gray pup over the black and white pups in both interactive and non-interactive runs; the effect was a bit more pronounced in interactive runs.

#### ***Identification with Wolves***

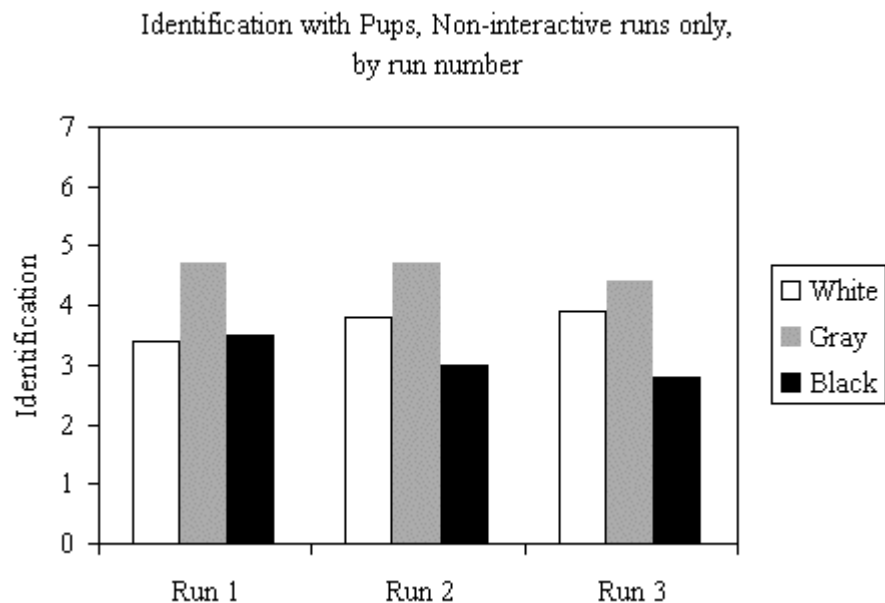
A related set of questions in the initial questionnaires pertained to how much people identified with the wolves. Comparing the 64 interactive runs with the 32 Non-interactive runs, subjects gave the white pup an average score of 3.5 in interactive runs and a 3.6 in Non-interactive runs ( $p = 0.6821$ ). These subjects gave the gray pup an average score of 5.1 in interactive runs and a 4.7 in Non-interactive runs ( $p = 0.2486$ ). They gave the black pup an average score of 3.5 in interactive runs and a 3.2 in Non-interactive runs ( $p = 0.2810$ ).

Looking at these data from the point of view of each individual pup, there is a clearly greater amount of identification with the gray pup in both interactive and non-interactive runs (see Figure 5-17). In the 64 interactive runs, the average score for identification with the gray pup was 5.1, versus 3.5 for either of the other pups ( $p = 0.0006$ ). In the 32 Non-interactive runs, the gray pup scored a 4.7 and the other pups averaged 3.4 ( $p = 0.0017$ ). Clearly, people identify more with the gray pup, even when they're just watching. Even so, the interactive runs showed a slightly more exaggerated effect.



**Figure 5-17: Cinematography matters more than interactivity in identification, too.**

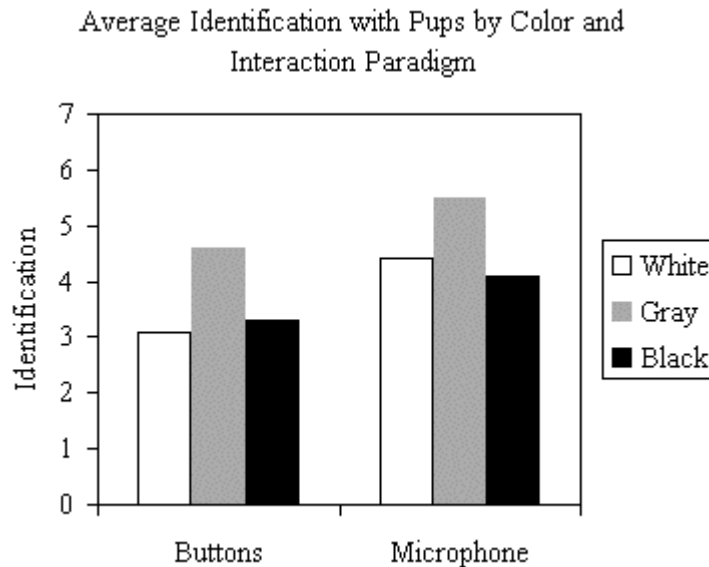
There was a strong disparity in identification between the gray pup and the others even in Non-interactive runs in the first position, reaffirming that subjects' identification with gray had little to do with any cross-over effects from other runs when the subject was directing that pup. Gray outsourced its siblings by 1.3 points on the seven point scale on Non-interactive runs in the first position, by 1.4 points on Non-interactive runs in the second position, and by 1.1 points on Non-interactive runs in the third position (see Figure 5-18). (Interestingly, gray outsourced its siblings by only 0.7 points during interactive runs in the first position. People identified with the gray pup in the first run *less* when they were controlling it.)



**Figure 5-18: People identify more with the gray pup in non-interactive runs, even when it's their first run.**

Among the eight subjects with Microphone runs, people appeared to identify more closely with all wolves in Microphone runs. They gave the white pup a score of 4.4 in the Microphone runs versus a score of 3.1 in Button runs ( $p = 0.1949$ ), the gray pup scores of 5.5 and 4.6 respectively ( $p = 0.2786$ ), and the black pup scores of 4.1 and 3.3 ( $p = 0.3823$ ). While the sample size was not large enough to reach statistical significance, subjects appear to identify more closely with all three virtual wolves when using the Microphone interface.

In both Button and Microphone runs for these eight subjects, there was more identification with the gray pup (see Figure 5-19). In the Button runs gray outsourced the other pups 4.6 to 3.2 ( $p = 0.1304$ ) and in the Microphone runs 5.5 to 4.3 ( $p = 0.0873$ ).



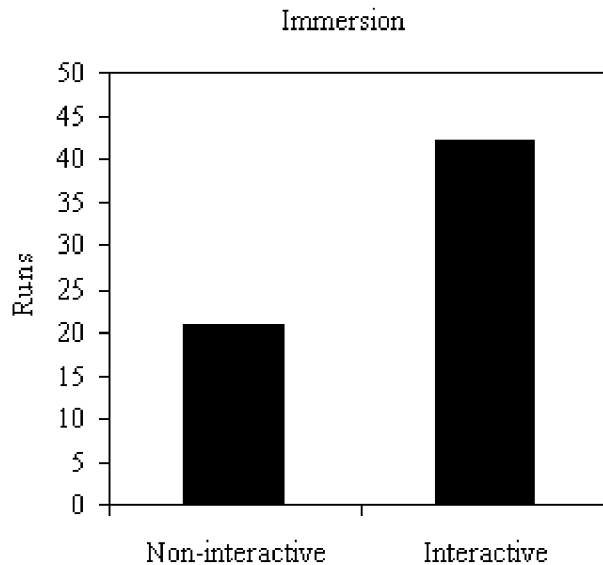
**Figure 5-19: Subjects identify more with all three pups when interacting with a microphone.**

To summarize, subjects identified with the gray pup more than with the black and white pups in both Non-interactive and interactive runs. This effect is more pronounced in the interactive runs. Within the interactive runs, people identified more closely with the wolves during the Microphone runs than in the Button runs. The identification with the gray pup was even more exaggerated than in the Button runs.

### ***Immersion***

The final area in which we compared the various interaction paradigms is how immersed people felt in each run. The questions to which subjects responded was: “In which of the runs did you feel more immersed in the experience?”

Interactive versions of the virtual wolves proved to be more immersing than Non-interactive runs. Of 64 responses, 42 chose the interactive run, 21 chose the Non-interactive run, and one tied (66.7%,  $p = 0.0056$ , see Figure 5-20). Of the eight subjects with both Button and Microphone runs, five felt the Microphone run was more immersing (63%,  $p = 0.3633$ ).



**Figure 5-20: Interactive experiences with the wolves are more immersing than non-interactive experiences.**

One 29-year-old male subject reported an element of the interaction that didn't work for him: "Prescribed relationships in button runs greatly reduces feeling of immersion. Should let us duke it out with other pups." Nevertheless, in order to perform the encoding/decoding experiments above, it was necessary to specify the relationships.

To summarize, while there was not enough data to distinguish between Button and Microphone runs, it was clear that interactive runs are significantly more immersing than non-interactive runs.

### **Discussion**

There are two main sets of phenomena that have become apparent from the above results.

First, interactivity appears to conflict with realism in the virtual wolves, and to decrease the clarity of relationships. Nevertheless, interactivity allows subjects to have more control and also creates a more immersing experience. One 44-year-old female subject spoke directly to this inverse relationship: "My comfort level with + memory of virtual wolves' behavior increased proportionately with the decrease in my direct control (responsibility) for their behavior." The interaction paradigm between participant and virtual wolf, with the pup being a separate entity in certain respects and yet directable by the participant, may have had a significant effect on the relationship between realism and immersion.

Second, people appear to like wolves more and identify more closely with them when they're the camera's object of attention, and even more when they can direct them too. Within the interactive runs, the Microphone interaction paradigm appears to increase the identification that people feel with the pups, although it doesn't make them like the pups more.

In the next chapter, I address the relationships among these topics more fully.

#### 5.2.3.4 *Additional Thoughts*

In addition to the various quantitative results described in the preceding pages, there is some interesting anecdotal evidence to be gleaned from the comments that subjects wrote at the end of each questionnaire, when asked for “any additional thoughts.”

##### **Pronouns and the Relationship between Subject and Pup**

In their comments, many subjects made reference to the gray pup. How, exactly, they made that reference, though, showed some variation. Many subjects (at least 8), referred to the gray pup as “I” or “me.” For example, one 28-year-old female subject described an interaction with another pup: “I wanted it to come after me so I could whine + be pitiful.” In the interaction, she became the gray pup (or it became her).

Several subjects (approximately 5) referred to the gray pup in the third person. For example, one 28-year old male subject said, of the second interactive run, “I had an easier time with this one, getting the gray pup to do what I wanted.” His terminology shows a clear distinction between the autonomous pup and his directorial control.

A few subjects showed ambiguity in the area of pronouns. A 30-year-old male subject started off talking about “my pup’s first response”, and how “my pup wasn’t focusing”, but in a later run shifted to the first-person, using phrases such as: “a pup and I growling at each other.” This subject’s progression from the third-person to the first-person suggests that perhaps disobedient pups are seen as separate, uncontrollable entities, whereas a pup who does exactly as it is instructed can become a vessel for the self. (Both of the subject’s interactive runs were Button-controlled, so interface was not a factor.)

This ambiguity in terms of the relationship between subject and pup is also evident in the following quote from a 32-year-old female subject: “Hard to tell dominance – as if I couldn’t get the white pup to dominate – it would see the grey dominate black + then submit to grey as well. But – if I were actually the grey – I would take advantage of this.” This passage suggests that the subject was put into conflict when asked to submit to another pup who clearly did not want to dominate her. If she and her pup were the same entity, she would have tried to dominate the other pup. It is interesting to consider the ramifications of game-design elements on the relationship being constructed between the player and the character.

The gender that people attributed to the pups was also interesting. Most of the subjects who used a third-person pronoun referred to the pups as males, despite the fact that the experimenter consciously and rigorously used the neuter pronoun “it” whenever mentioning the pups. Of the eight respondents who used a third-person pronoun, five used “he” or “him”, two used “it”, and only one used “her.” In fact, the 28-year-old female subject who used the pronoun “her” first wrote “him”, and then crossed it out. “My wolf wanted to sniff the ground too often when I wanted ~~him~~ her to be dominating the black pup.” That subject, at least, had a strong opinion about the gender of the pup. In a further study, it would be interesting to determine if the perceived gender of the pups changed if the interaction featured behaviors beyond dominance and submission (e.g., play, social grooming, etc.).

While there is little quantitative evidence to be had from these comments, they shed some light on the complex way in which people identify with, control, and become characters in interactive media.

### **Telling Stories**

Subjects sometimes showed an emotional attachment to the pups in their end comments, and were beginning to construct narratives around them. A 21-year-old female subject wrote: “I felt sorry for the white one.” This is an example of the installation *working*. She was drawn in, past the graphics, past the interface, past the experimental setting, and had an emotional response to at least one of the pups.

A 14-year-old male subject in the pilot study wrote a thorough description of his interaction, including the following: “I didn’t like the white wolf because he was so power-hungry, on the other hand, I didn’t really like the black wolf because he gave in to the white one so easily.” For some subjects, the wolves began to have motivations, agendas, and life-histories.

### **Enthusiasm and Learning**

Many subjects comments included positive feedback. For example, one 18-year-old female offered: “Very cool simulations!” In addition to these positive but non-committal comments, several subjects connected the wolves to learning about social behavior. A 30-year old male wrote: “I want to see more of that video now!” A 28-year-old female wrote: “I think I could be better at being submissive after watching this. It is more helpful than Nat[ional] Geo[graphic] Video.” Making an interactive installation that could help people learn about the social behavior of the wild wolf has been one of the central goals of the AlphaWolf implementation.

Several subjects drew connections with experiences in the real world, or with their own behavior. A 21-year-old female wrote: “It was also interesting to me to see how the dominance routines are so similar to what I watch kittens do.” A 23-year-old male wrote: “I was thinking about my own behavior (as the grey pup) and noticed I tended to dominate white whenever I saw them, but waited for a dominance display before submitting to black.” Getting people to think about social behavior, and in particular about their own behavior, is one of the deep goals of this entire project, and will be addressed more fully in the Applications section below.

### **Expressiveness**

Several subjects commented on the expressiveness of the characters. A 22-year-old female wrote: “Beautiful animation. Incorporation of the ears & tail in addition to larger body motions helpful in reading submissive/aggressive postures.” A 29-year old male put it more bluntly: “The white pup submitted to me very well. Wuss.”

One 24-year-old male subject even implied that the expressiveness was too extreme: “Submissive wolves seemed too cowardly.” The fact that the submissive wolves were seen as “too cowardly” means that we hit the mark in making very expressive characters.

### **Problems**

Not all of the feedback was positive. There were several areas that more than one subject mentioned – interpenetration, sound effects, and camera work.

#### ***Close Contact vs. Interpenetration***

Sometimes the virtual wolves, despite elaborate collision-avoidance systems, interpenetrate, for example by walking through each other. A number of people mentioned the interpenetration – often humorously. For example, one subject in the pilot study commented that: “[t]he dominance relationship of graphical superimposition is unclear. ☺” To people interacting with the wolves, the interpenetration is the most obvious technical problem.



Nevertheless, several people mentioned the lack of neck-biting as a gesture of dominance, which featured prominently in the real wolf video. It is exactly this kind of close contact that is not possible using our current system, and which is currently an unsolved problem in computer graphics. We all agree that it would be incredibly powerful to have wolves who could be so aggressive that they grabbed another wolf by the muzzle and pinned it to the ground. Currently, though, any implementation approaching this kind of behavior is fraught with heads-inside-of-heads and other gruesome violations of the suspension of disbelief. To show dominance and submission with a significantly higher level of realism, we need close-contact motor-control.

### ***Sound Effects***

A few people mentioned the sound effects negatively. One 21-year-old male mentioned that the “whining seemed a bit too ‘Star-Trek’-y.” In the course of 12 minutes of interaction, many of the sound samples were repeated; some of the samples are sufficiently recognizable that this repetition might be quite obvious. Sound design for synthetic characters, in particular for sound-intensive topics like wolf dominance and submission, is a sorely under-addressed research topic.

### ***Camera***

The third topic that gave several subjects trouble was the camera work. Interactive cinematography for virtual environments is a challenging research topic, and has been addressed previously by the author [Tomlinson 2000b]. Despite the wolves having a state-of-the-art camera system, one subject reported feeling a little dizzy as a result of the camera’s motion, and another complained of the camera not showing key interactions among the pups.

Nevertheless, subjects were very successful at encoding and decoding the wolves’ relationships, tasks made possible by the camera system. Also, as became clear in the sections on “Liking of Wolves” and “Identification with Wolves”, the camera had a significant effect on the way people felt about the different wolves. By focusing on the gray pup, the camera caused subjects to identify with the gray pup and like it better than the other two. Therefore, the camera succeeded in its central purpose – to show people what they needed to see – and also had the side-benefit of helping build a unique relationship between the interactor and the pup being directed.

### 5.3 Simulations

The second set of evaluations that I performed on the AlphaWolf system involved choosing a variety of real-world problems and seeing if the wolves' social relationships made it easy to simulate them. These simulations all involved a pack of six virtual wolves consuming a virtual carcass. This scenario mimics a situation seen in the wild when limited resources force individuals to compete for food. As David Mech mentions, "the most practical effect of social dominance is to allow the dominant individual the choice of to whom to allot food." [Mech 1999]

In addition, the problem of a group of wolves feeding on a single carcass resembles a problem that roboticists face in trying to get multiple robots to recharge at a single power source as efficiently as possible [Michaud 2001]. "As the size of the group grows, ... interference increases, causing the decline in global performance, and presenting an impetus for the use of social rules" [Mataric 1995]. The larger goal of these synthetic experiments is to provide support for the potential application of the AlphaWolf mechanism to multi-robot systems and other groups of interacting computational entities.

The several goals of these simulations were:

**Goal 1:** To create a group of autonomous entities that form a distributed linear dominance hierarchy, and verify that this hierarchy is caused by the AlphaWolf mechanism.

**Goal 2:** To examine the effect of this dominance hierarchy on resource exploitation, comparing it to other kinds of social structures.

**Goal 3:** To make a multi-agent system in which an external authority can restructure the distributed linear hierarchy in a controlled way by altering the behavior of a single individual.

Each of these specific goals support one central hypothesis:

**Central hypothesis:** The AlphaWolf mechanism is a simple, robust, social relationship mechanism that can be used as the basis for a variety of tasks involving multiple interacting computational entities.

The process of achieving the three goals listed above began with an assortment of packs of virtual wolves in simulated environments. The wolves and their worlds were based on the AlphaWolf installation that was shown at SIGGRAPH and used in the human-subject study, but had several differences. To interact with the carcass, wolves had an augmented perception system and an additional set of feeding behaviors. In addition, to increase the speed with which the wolves formed relationships, they did not spend as much of their time wandering around as previous wolves had. The distinctions between the various versions of the wolves is addressed more fully in the Implementation chapter of this thesis.

Each run of virtual wolf pups occurred in near-real-time, each taking several hours to establish the 30 different relationships (each of 6 wolves having a CSEM of the 5 others) and view the effect of those relationships on the task of carcass-exploitation. The pups were colored to keep them straight – gray, black, white, red, green, and blue. In the various graphs below, pups are shown as their color.

### 5.3.1 AlphaWolves Form Hierarchies

Real wolves in captivity form dominance hierarchies [Fox 1971]. In wild wolves, the family structure also shows somewhat hierarchical arrangements, with the parents controlling food access among their offspring [Mech 1999]. At first glance, dominance hierarchies seem unfair to lower ranking wolves. However, it is far more efficient than having lengthy fights every time two individuals want the same thing. Ultimately, if the wolf who *would have lost* the fight is willing to capitulate, the result (in terms of who gets to eat first) is the same, and both wolves are spared the significant costs (in terms of energy, time and potential injury) associated with having an actual fight. Dominance hierarchies are a technology for efficiently arbitrating among individuals who are able to remember previous interactions with other individuals.

To show that virtual wolves with the AlphaWolf mechanism tend to form dominance hierarchies, I placed 6 identical naïve wolf pups in a virtual world with a carcass on which only one wolf could feed at a time. Each pup began with no emotional memories, and the same behavioral repertoire. Over the course of approximately 5000 virtual seconds (approximately 4 hours of real time), the wolves differentiated from each other by means of their social relationship mechanism, taking on an assortment of dominance roles.

The way in which I calculate the dominance value of a wolf who has a number of different relationships with its pack mates is to take the average of the dominance value of each relationship. Through this method, dominance hierarchies become plainly visible, with more dominant wolves appearing higher on the graph. Other techniques for determining a linear hierarchy from a group of interacting individuals have been proposed (e.g., [de Vries 1998]), but simply showing the average internal state suffices for the purpose here.

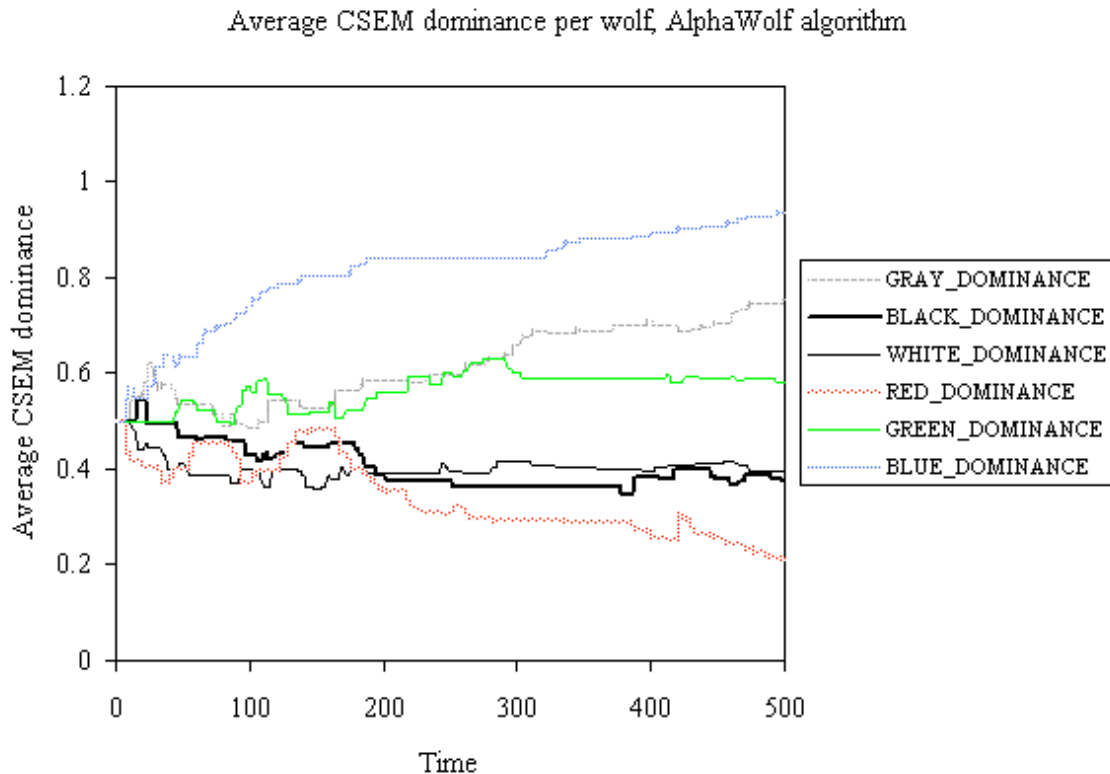
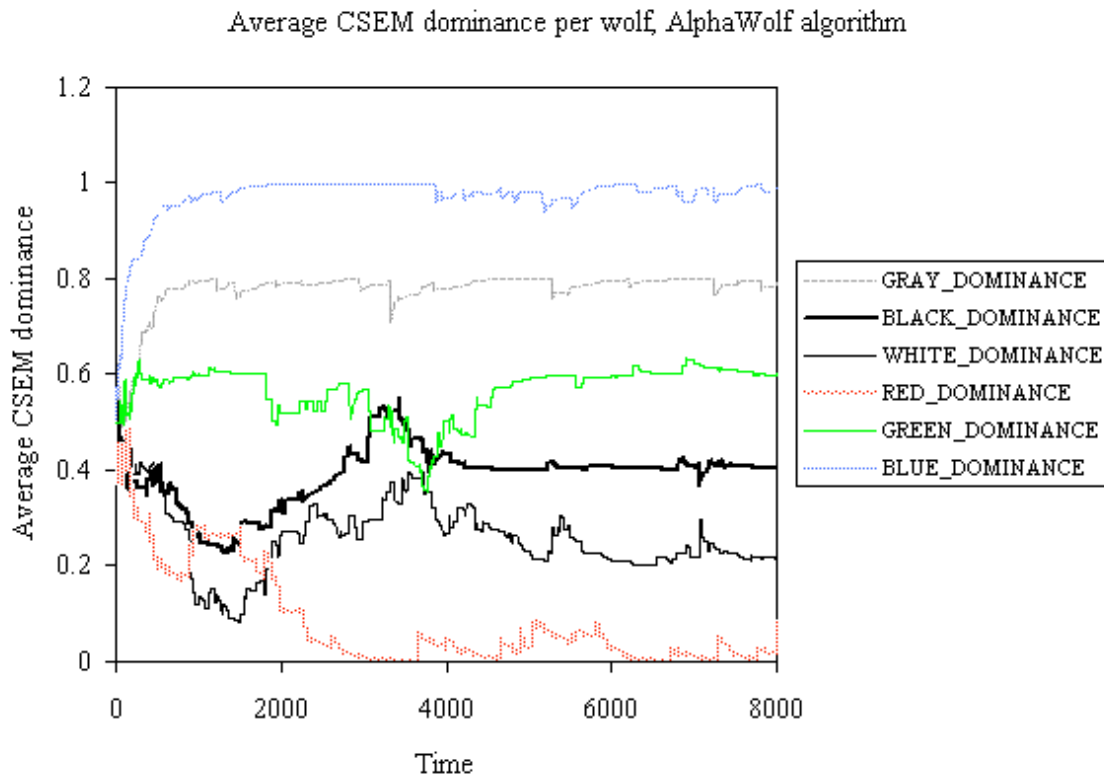


Figure 5-21: The first 500 virtual seconds in a newly created pack.

In the early stages of their time together, the pups often switched dominance roles (see Figure 5-21). Since none of the pups were very confident in their relationships, they were easily changed. Real wolf pups, too, being their lives with very fluid relationships and develop a hierarchy between eight and twelve weeks of age [Fox 1971].

Later in the pack's life, the dominance roles become more extreme, and an evident dominance hierarchy begins to emerge (see Figure 5-22). One pup consistently wins all of its interactions, becoming the “alpha” pup (in this run, the blue pup became the alpha individual). Another pup always loses, becoming the “omega” pup (in this run, red is at the bottom of the hierarchy). In between these two extremes, the other four wolves work out relationships, too, taking on specific, consistent roles with respect to each other pup. The emergent structure is the result of a group of local, dyadic interactions, rather than of some centralized authority.



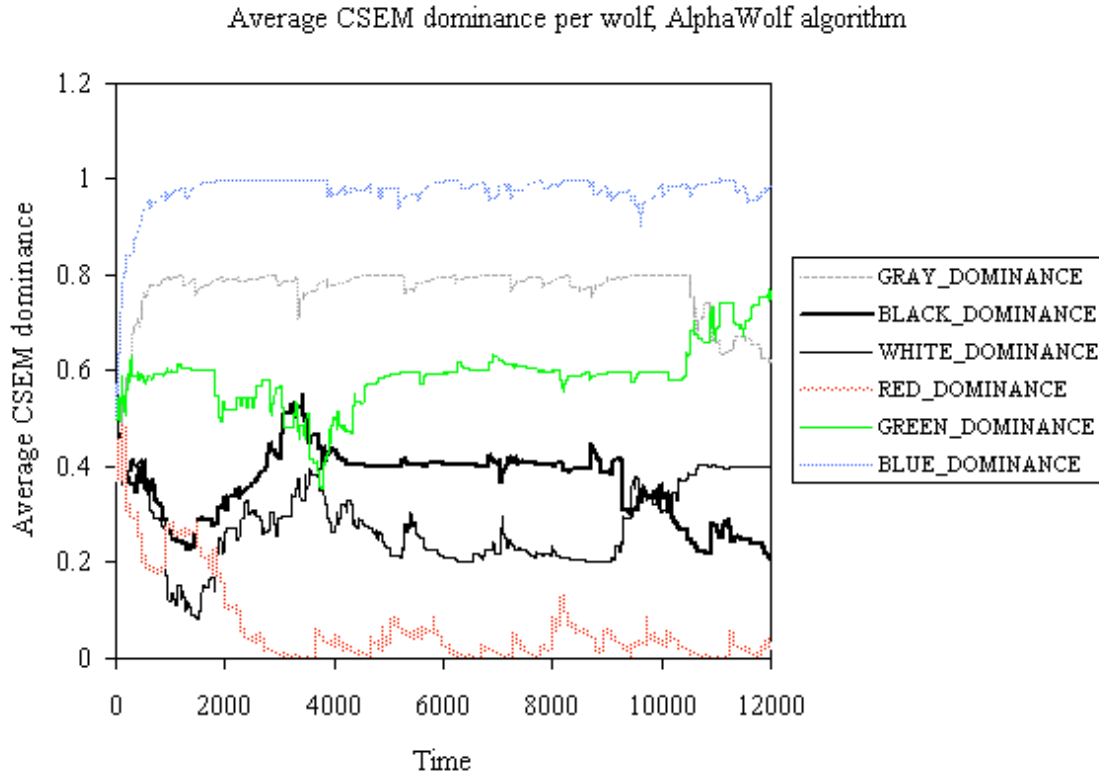
**Figure 5-22: The same pack as in the previous figure, now 8000 seconds old. The dominance relationships have largely stabilized.**

In these graphs, it is noticeable when two wolves have an interaction that is inconsistent with their dominance relationship. At approximately time = 7000, for example, white and black appear to have had a bit of an altercation. In real wolves, this kind of “testing” goes on frequently [Klinghammer 1985], with most fluctuations being ironed out by the stabilizing force of the relationship mechanism.

Nevertheless, there are occasional cross-overs in the hierarchy, when two pups (usually adjacent in rank) flip their dominance relationships. Examples of cross-overs can be seen in Figure 5-23, at time = 9500 between black and white, and at time 11000 between gray and green.

In the AlphaWolf system, cross-overs appear to occur more frequently among middle-ranked wolves, who are more prone to exhibiting both dominant behaviors (to the lower-ranking wolves)

and submissive behaviors (to the more dominant wolves). Interestingly, this tendency mirrors the structure of captive wolf packs, where there is often a clear alpha individual for each gender (separating hierarchies by gender will be addressed in the Discussion section below), a clear omega individual, and a bit more fluid set of relationships among mid-ranking wolves [Klinghammer 1985].



**Figure 5-23: The same pack, with crossing-over among middle-ranking wolves.**

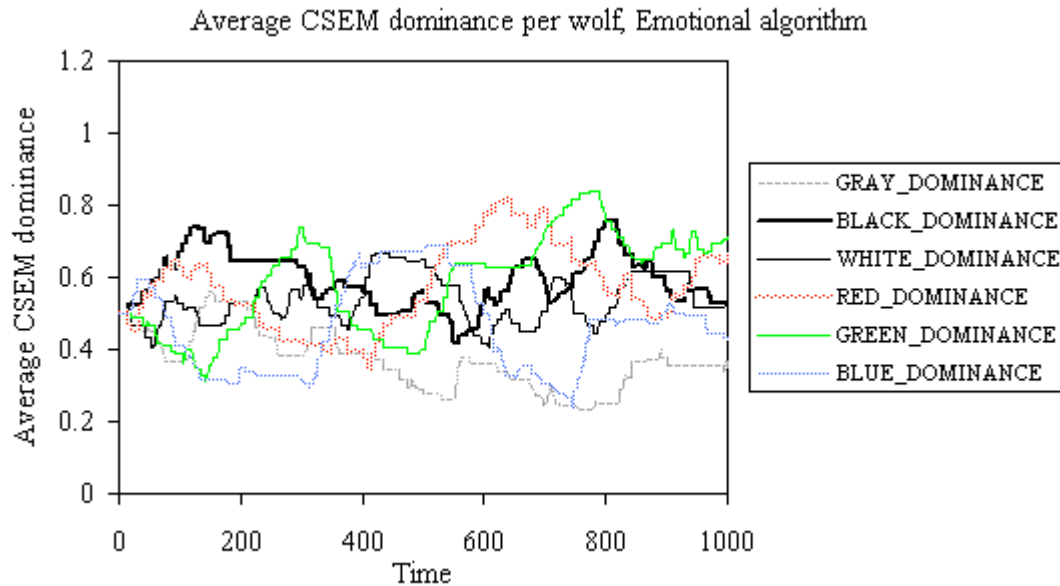
The fact that the AlphaWolf algorithm allows a group of naïve individuals to work out a hierarchy suggests that it would also be possible to introduce new individuals to a pre-existing hierarchy. Just as new pups are born once a year among wolves [Mech 1998], new autonomous agents could find their place in the existing social structure.

#### **5.3.1.1 Other Algorithms**

To show the differences in social structure that are possible among a group of virtual entities, I ran three other packs, one with each of the alternate algorithms that had been implemented for the human-user studies. These three algorithms are:

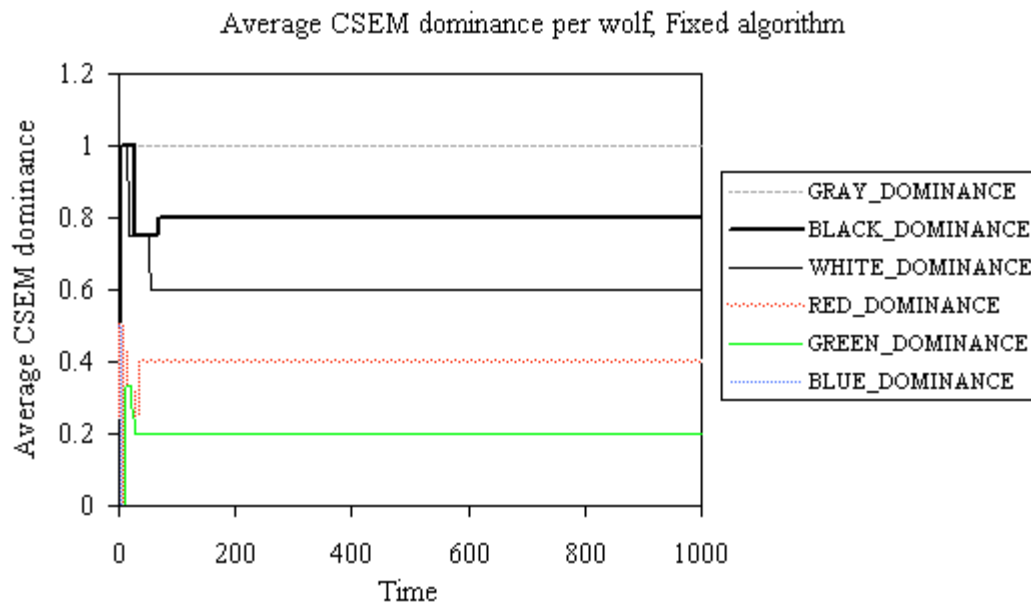
- Emotional – Each wolf has a dynamic emotional state that is affected by its interactions, and which affects its action selection mechanism. This algorithm is identical to the AlphaWolf mechanism except that the CSEMs are prevented from affecting the current emotional state.
- Fixed – Each wolf is loaded with a pre-defined set of CSEMs, that together make up a linear hierarchy. The wolves' interactions do not change these CSEMs, so their ranks are stable and fixed.
- Random – Each wolf chooses a random value for its CSEM each time it starts to interact with another wolf.

The graphs produced by the different packs give a clear flavor for the kind of social order (or lack thereof) that they cause. Each of the graphs below shows the first 1000 virtual seconds of a pack with one of the relationship mechanisms. In each case, the graph continues in the same fashion for as long as I have collected data (8000-10000 ticks).



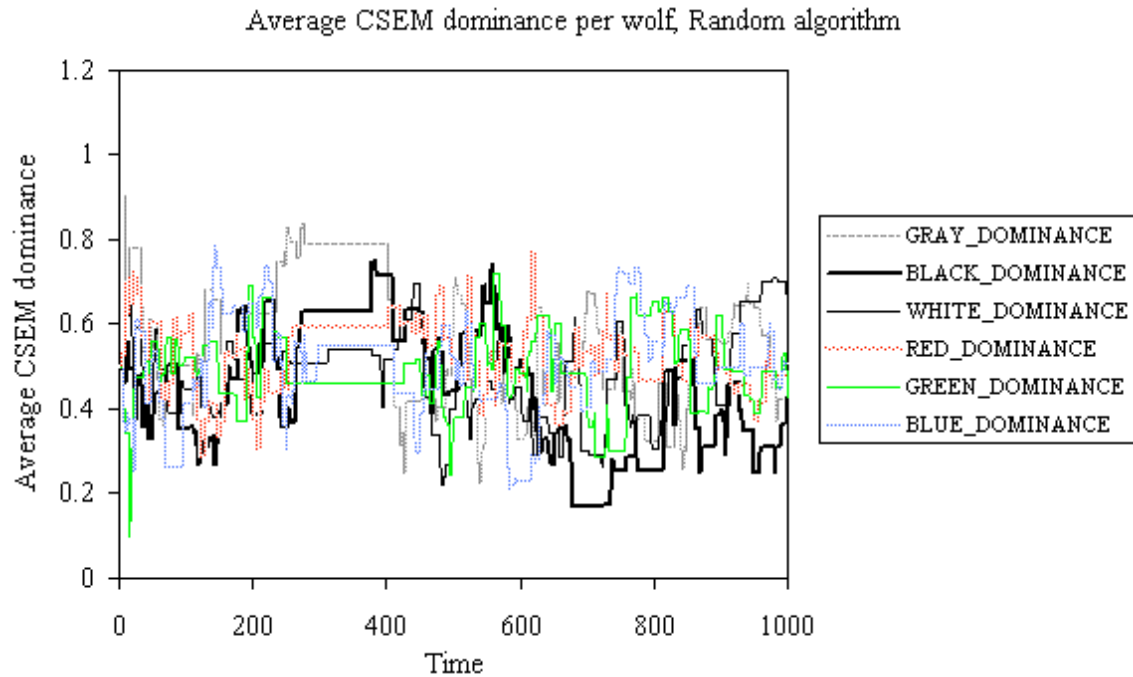
**Figure 5-24: The first 1000 seconds of a pack featuring the Emotional algorithm.**

The Emotional algorithm (see Figure 5-24) causes fairly rapid turn-over among the pups. Nevertheless, there is a certain amount of persistence in each pup's emotional state. The graph is based on the emotional memories that the pup forms. Even though the emotional memories do not affect the actual emotional state of the pups, they do provide a reasonable metric by which to visualize the place of the wolves in the social order of the pack.



**Figure 5-25: The first 1000 seconds of a pack with the Fixed algorithm.**

The Fixed algorithm (see Figure 5-25) creates a strict hierarchy that does not change based on the pups' interactions. Each wolf rapidly assumes its place in the pack (since each fixed CSEM kicks in the first time two wolves meet) and thereafter behaves appropriately given its relationships with its pack mates.



**Figure 5-26: The first 1000 seconds of a pack with Random social relationships.**

The Random algorithm (see Figure 5-26) results in constant rapid change. By comparison, the Emotional algorithm looks positively leisurely.

Each of these other algorithms demonstrates alternative social possibilities to the one created by the AlphaWolf mechanism. The Fixed algorithm offers a strong hierarchy, but one which is inflexible and unchangeable. The Emotional algorithm offers rapid adaptability but has persistence on only a short time scale. The Random algorithm provides a view into the chaos that is possible if some social structure is not enabled. The social structure among wolves with the AlphaWolf mechanism shows the most similarities to the social structures of real wolves, with consistency but flexibility.

The above graphs confirm that I was able to achieve **Goal 1**, creating a virtual pack with a distributed, essentially linear, dominance hierarchy that is clearly caused by the AlphaWolf mechanism.

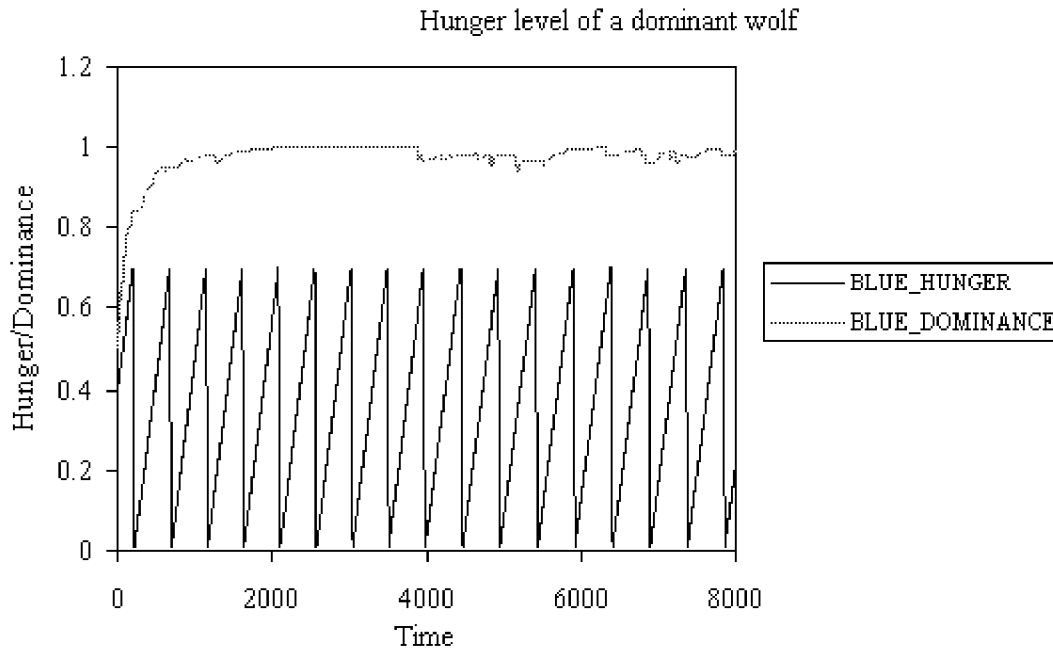
### 5.3.2 Hierarchies Create an Imbalance in Resource Distribution

In the second section of these simulations, I explore the effect of the various social structures on resource exploitation. To display this effect, I recorded the hunger of each wolf once per virtual second during the four runs above, one with each algorithm (AlphaWolf, Emotional, Fixed, Random).

In real wolves, there appears to be a small area around a hungry wolf's mouth where dominance does not apply; with respect to food, ownership is nine-tenths of the law [Mech 1970]. Nevertheless, higher ranking wolves appear to have priority when it comes to securing food

[Mech 1999]. To mimic this phenomenon in simulation, I caused wolves to voluntarily relinquish the carcass when a wolf whom they believed to be more dominant came within a certain distance. This voluntary departure could help prevent physical interference among robots, since the submissive entity will give way for the dominant entity to feed.

The result of this behavior is clear – that lower-ranking wolves tend to be hungrier than higher ranking wolves. Their feeding gets interrupted frequently by other members of the pack coming to feed. In Figure 5-27, we see the hunger level of a dominant wolf, matched with a graph of its dominance. This wolf approaches the carcass when it is hungry (hunger = 0.7), feeds until it is full (hunger = 0.05), and then goes about its business until it is again hungry.



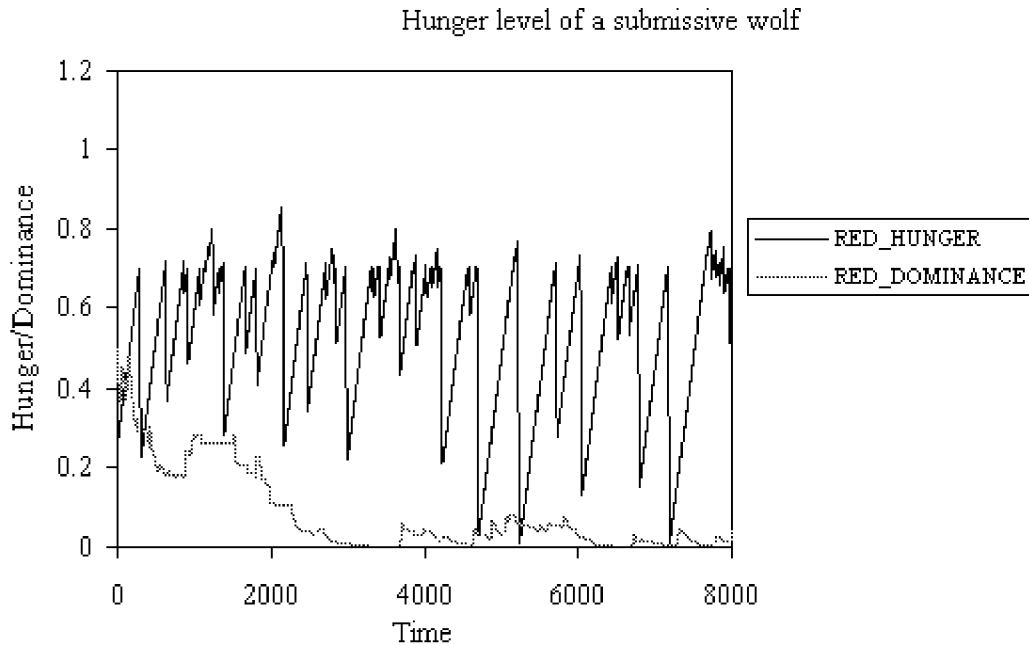
**Figure 5-27: A dominant wolf can feed whenever it wants (here, when hunger reaches 0.7).**

On the other hand, a submissive wolf from the same pack shows a very different hunger profile (see Figure 5-28). It spends much more of its time lurking around the carcass waiting for a bite to eat.

While this pack has enough to eat (i.e., the hunger growth rate and the feeding rate are set such that the entire pack can be easily fed) it is interesting to consider what would happen if food were made more scarce. The dominant wolves would still eat their fill, but lower ranking wolves would have a harder and harder time keeping their hunger at bay.

This imbalance in resource distribution favoring dominant individuals is exactly what happens in real wolf packs, where food scarcity causes pack size to decrease [Mech 1998]. Lower ranking wolves go hungry, and eventually split off to form their own packs where they will be able to find and consume their own food.



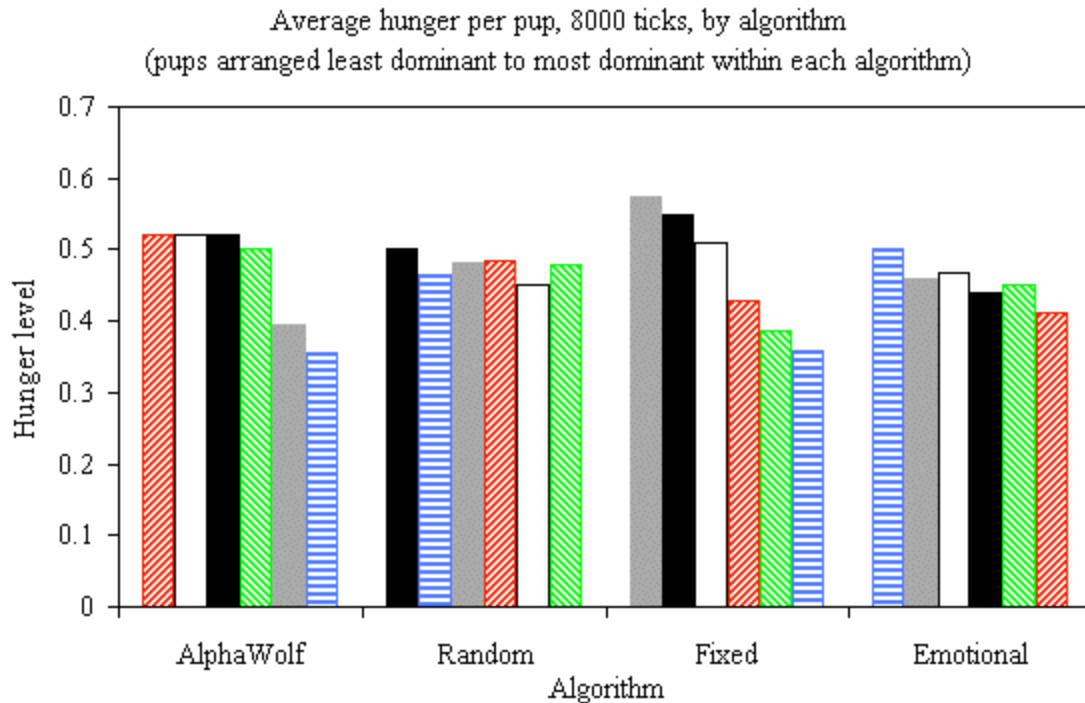


**Figure 5-28: A submissive wolf often gets interrupted when feeding, resulting in a much more jagged hunger profile.**

In addition, an imbalance in resource distribution might be desired behavior in a multi-robot system. Rather than having a homogeneous group of robots, there might be important differences. Imagine, for example a colony of robots on Mars, in which some robots are responsible for research and others are responsible for maintaining the base station. Causing the research robots to be “dominant” over the cleaning robots would allow them priority access to fuel, thereby operating as efficiently as possible, but at the expense of reduced efficiency in the cleaning robots.

Returning to the four different social relationship algorithms, we can see clearly the different degrees that the imbalance in resource distribution takes (see Figure 5-29). The Fixed algorithm shows a distinct spread of hunger levels (0.36 to 0.57) arranged by dominance, whereas the Random algorithm shows almost no hunger spread (0.45 to 0.50). The AlphaWolf pack shows a relatively large spread (0.36 to 0.52) and the Emotional pack less so (0.41 to 0.50). These figures suggest that, while AlphaWolf doesn’t work quite as well as a Fixed, nonnegotiable hierarchy, it nevertheless shows a strong imbalance of resource distribution. In addition, I expect that some of the dominance ambiguity in the AlphaWolf figures is the result of the initial period when the dominance hierarchies are still being learned, and would be less and less of a factor as time went on.

The variation in hunger-spread is caused to a large extent by the variation in dominance. While the Fixed mechanism shows a difference of 1.0 between the average dominance of the highest ranking individual and the lowest, AlphaWolf shows a difference of 0.9, Emotional shows a difference of 0.2, and Random shows a difference of only 0.05. Since hunger is directly correlated with rank, a strong dominance hierarchy results in a correspondingly strong hunger hierarchy.



**Figure 5-29: The Fixed algorithms creates a pack with strong polarization of hunger levels. AlphaWolf is next, followed by Emotional. The Random algorithm shows very little polarization of hunger among the pack members.**

Despite the differences in hunger among the wolves in these four packs, none of the wolves in any of the packs ever starved to death (i.e. hunger reached 1.0). Nevertheless, reducing the quality of the available food did cause the lower ranking wolves to hit the starvation level. In real wolf packs, though, low-ranking individuals would disperse before they starved [Mech 1998].

Comparing the dominance-based resource allocation of the AlphaWolf system to an altruism based system in which the a wolf will yield a food source to a hungrier wolf, the altruism based system would show less of a disparity in hunger. While an altruism based system might be more appropriate for a group of uniform robots, the dominance based system would allow certain robots to exhibit a higher performance level to others.

The studies above address **Goal 2**, examining the effect of the various social structures on resource exploitation. To summarize, dominance hierarchies allow an imbalance in resource distribution, and the AlphaWolf mechanism makes hierarchies that show many of the characteristics of real wolf social groups.

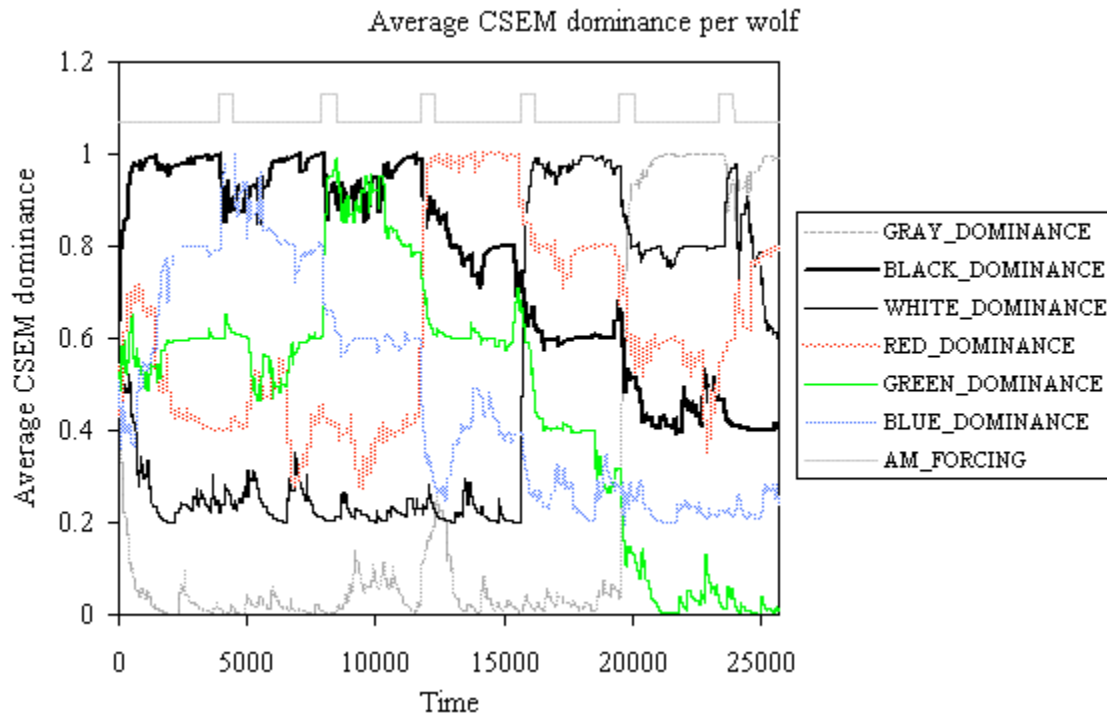
### 5.3.3 AlphaWolf Hierarchies Are Directable

The third goal of these simulations was to make a directable hierarchy of agents – a multi-agent system in which an external force can restructure the distributed linear hierarchy in a controlled way by altering the behavior of a single individual.

In nature, external forces frequently intervene to restructure social hierarchies. A dominant wolf is sometimes killed in a hunt or an accident, and its breeding partners may take a new mate [Mech 1998].

In a multi-agent or multi-robot system, too, it might be appropriate to have a hierarchy that can be intentionally restructured by an external force. Returning to an example with research robots and cleaning robots from the previous section, under normal circumstances, the research robots might deserve first access to power, so that they can go on longer excursions without refueling. However, if there has just been a dust storm and the need for cleaning is hampering the productivity of the whole colony, then it might be appropriate to shift priority access to the cleaning staff. The now-dominant cleaning robots are now more efficient, at the price of “hungrier” research robots.

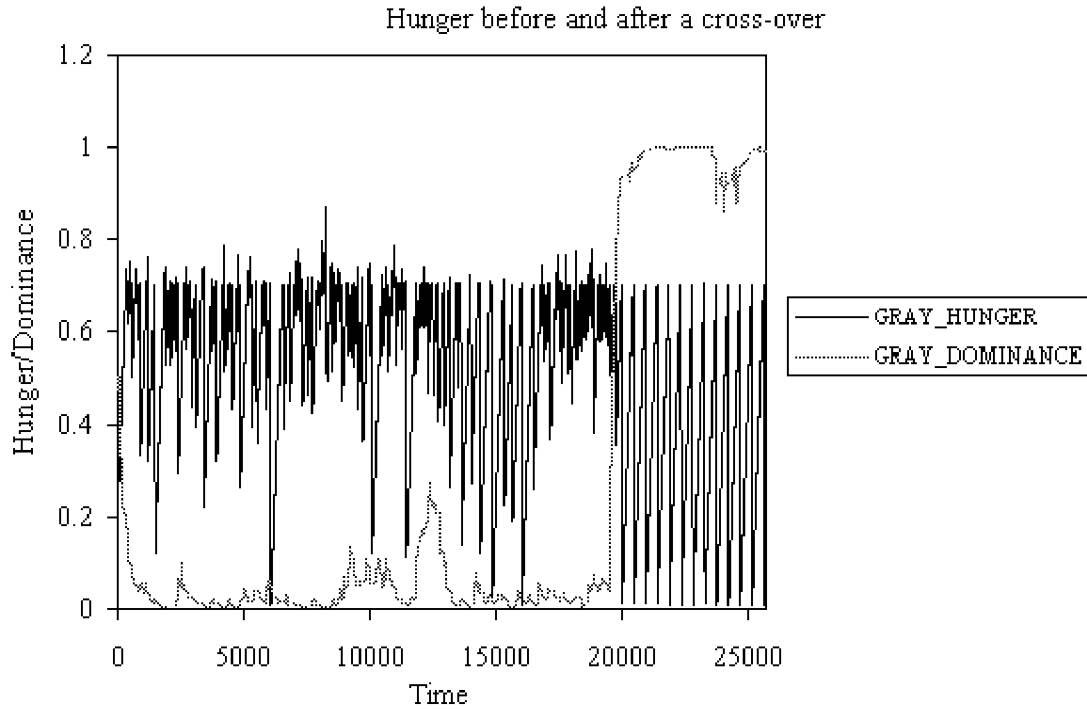
To test the tolerance of AlphaWolf hierarchies to external direction, I created a pack of wolves who were periodically subjected to an external influence that caused one wolf to remain dominant through all its interactions (see Figure 5-30). Every 4000 virtual seconds, a certain wolf was chosen and forced to be dominant for 600 virtual seconds. Each time a wolf was chosen it was the next lower-ranking wolf, beginning with the second-ranking wolf. For example, at  $t = 4000$ , the second-ranking wolf was forced dominant until  $t = 4600$ . Subsequently, the third-ranking wolf was forced dominant from  $t = 8000$  to  $t = 8600$ . Each rank was promoted in turn, regardless of which wolf currently held that rank. The square wave at the top of the graph shows the exact points when external direction was turned on and off.



**Figure 5-30: A pack of entities with the AlphaWolf mechanism are directable. Every 4000 seconds, a wolf is chosen for “promotion”, and is forced to be dominant in all its interactions for 600 seconds. The effect is most evident among the promotions of lower-ranking wolves (Red at  $t = 12000$ , White at  $t = 16000$ , Gray at  $t = 20000$ ).**

Interestingly, it appears easier for lower-ranking wolves to be promoted to the top than for middle-ranking wolves. In this graph, the three “coups” staged involving a 2<sup>nd</sup> or 3<sup>rd</sup> ranked wolf ( $t = 4000$ ,  $t = 8000$ ,  $t = 24000$ ) failed to promote that wolf, while the three coups involving a 4<sup>th</sup>, 5<sup>th</sup>, or 6<sup>th</sup> ranked wolf ( $t = 12000$ ,  $t = 16000$ ,  $t = 20000$ ) all succeeded. This phenomenon probably occurs because the more striking alterations of the hierarchy result in periods of

increased instability in the pack, thereby making it easier for the transplanted wolf to hold its new position.



**Figure 5-31: When a wolf is promoted, its hunger profile becomes that of a dominant wolf.**

The change of dominance role has a corresponding effect on the hunger profile of each wolf, too. Figure 5-31, for example, shows the hunger graph of a wolf who was the omega individual for most of the first 20000 virtual seconds and was then promoted to alpha status. After the promotion, its hunger graph demonstrates its dominance in the realm of feeding as well as simple social rank.

These simulations show that the AlphaWolf mechanism is amenable to external direction, thus satisfying **Goal 3**, to make a multi-agent system in which an external authority can restructure the distributed linear hierarchy in a controlled way by altering the behavior of a single individual.

### 5.3.4 Discussion

The fact that the three goals in the above sections were attainable supports the **central hypothesis** of these simulations, that the AlphaWolf mechanism is a simple, robust social relationship mechanism that can be used as the basis for a variety of tasks involving multiple interacting computational entities.

In the process of implementing these simulations, several points came to light that bear discussion. Having visual feedback from the wolves, for example, was invaluable in the process of debugging their social interactions. Simply watching a set of changing numbers would have made it nearly impossible to grasp the social interactions of this group of autonomous entities. Being able to see the gray wolf groveling at all comers, however, is crystal clear. While the runs took longer in a graphical world, being able to see the wolves' behaviors made it possible to accomplish the above goals with speed.

Several real-world phenomena came up that I did not attempt to implement. First, wolves exhibit separate hierarchies for males and females (at least among captive wolves) [Fox 1971]. These hierarchies nevertheless interact in various ways. This creation of separate hierarchies and yet interaction between them is beyond the current implementations that I have done. However, it too might be useful for multi-agent systems, if the agents were competing for several different resources. For example, perhaps the overlapping hierarchies result from competition for two real-world resources – feeding opportunities and mating opportunities. In terms of getting food, gender is irrelevant, so there might be competition among all members; in terms of mating opportunities, one male and one female must simultaneously “win.”

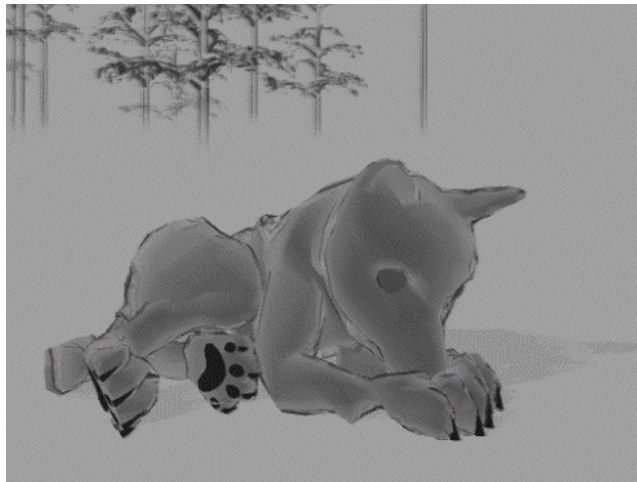
Another interesting real-world phenomenon that is not currently possible in the AlphaWolf system is that the alpha pair of wolves in a real pack give the youngest pups the right to feed before the middle-ranking members of the pack [Mech 1999]. This, too, might be behavior that would benefit a multi-robot system. For example, a certain robot performing a critical task might need a support staff of other robots who are only important because of their association with the primary robot. Giving the central robot the ability to endow its support staff with dominance might be similarly advantageous to the real-world case, in which parents insure the survival of their pups.

A third phenomena in real wolves that is not captured by the current model is the tendency of a deposed alpha wolf to drop drastically in rank, often becoming the omega wolf (also known as the “scape-goat”). In the current model, deposed alphas often drop only one rank.

I believe that the AlphaWolf mechanism could form the basis for a system that is capable of exhibiting these varied behaviors; I will discuss potential ways to implement each in the Chapter 7 below.

## 5.4 The AlphaWolf Installation

A third way in which we evaluated our mechanism was by observing the people who interacted with the AlphaWolf installation at SIGGRAPH. While we were unable to collect statistical information, the great majority of the 500-1000 people who interacted with the installation while visiting the Emerging Technologies program appeared to become immersed in their interactions with their pups.



**Figure 5-32: Each pup is asleep at the beginning of the interaction.**

Many participants were a little uneasy at first about the prospect of howling into a microphone in front of a crowd. The initial waking-up of the pups (see Figure 4) put users at ease with the relatively novel experience of interacting with a computational entity by means of a 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. The user *is not* the pup – the pup is asleep, the user awake. Nevertheless, it is *his* pup. Throughout the installation, the distinction between user and pup is clear, but there is still a strong empathetic association between the two. This association is focused by an intimate camera angle, a snoring sound effect, and a funny little ear-flip action that the sleeping pup does when it 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.

The immersion that people felt when interacting with their pups demonstrates that they experienced a sense of control over their pups. Participants were far more animated than they would have been watching a movie. Nevertheless, people often become quite animated when playing Doom and other “twitch games” without feeling that they are interacting with personality-rich characters. In the following paragraphs, we make several points that we feel shows that participants acknowledged the “character” of the wolves.

People often empathized with their pups, getting excited when the pups met new pack mates or “feeling bad” when somebody else dominated them. People mirroring the affect of their pups is a case of *social referencing* [Siegel 1999]. This suggests that those people were implicitly treating the wolves as social partners.

During and after their interaction with AlphaWolf, many people told stories about their pups to their companions. These stories usually focused on the relationships that their pup had developed with the other wolves. People’s ability and desire to tell these stories suggests that our wolves’ relationships were both clear enough for people to interpret and interesting enough for them to care.

These points, while admittedly unquantified, occurred frequently enough over the course of the SIGGRAPH run to be remarkable. We believe that these phenomena, combined with the excitement that people clearly felt while howling, growling, whining or barking for their wolves, demonstrate that people felt they were controlling characters who nevertheless maintained their own personalities.

#### **5.4.1.1 Press**

In addition, AlphaWolf received quite a bit of good press. While it’s hard to tell whether the press resulted from the social relationships, from the graphics renderer, from the Howl Interface, or just from the whole installation as a package, the press nevertheless supports the claim that people liked the installation. AlphaWolf made people talk to each other. Talking builds social relationships.

Here are a few excerpts:

“For a look at the truly experimental, convention attendees browsed displays of student projects, research-lab prototypes and industry tests in a separate exhibit called Emerging Technologies, to the sounds of howls, growls and whimpers from the Massachusetts Institute of Technology Media Laboratory’s Alpha Wolf project.

Participants could encourage their wolf to interact with other wolves in a computer-simulated pack, exploring the social behavior of the animals and getting a glimpse of the potential for autonomous virtual creatures.”

**Wall Street Journal Online, 8/21/01, “Growls and howls hit graphics show” [Grenier 2001]**

“MIT Media Lab probably wins the ‘way out’ interface award for its exhibit that ‘allows participants to interact socially with a pack of autonomous wolves’ by howling into a voice interface.”

**Wired News, 8/11/01, “Stretching it out at SIGGRAPH”, [Stroud 2001]**

“...one of the most beautiful pieces in the Art Gallery was Alpha Wolf, where people play the part of a wolf cub in a virtual pack, howling, growling or whimpering into microphone to influence their cub. Instead of trying to create photo realistic wolves, the Alpha Wolf team instead choose a style reminiscent of Chinese watercolor painting, and the results were much more evocative. Alpha Wolf was one of a several projects that tried to get rid of the mouse.”

**Interaction by Design, 8/17/01, “SIGGRAPH 2001 - is experimentation really dangerous?” [Olsen 2001]**

“The most enjoyable elements, and unique to SIGGRAPH, were the Emerging Technologies and the Art Gallery. These two areas showcased some truly interesting, if not bizarre creations - from liquid metal to wolf cub AI and magical wands...”

**International Game Developers Association, “SIGGRAPH 2001 Summary Report” [Della Rocca 2001]**

“If there is a current running through SIGGRAPH, from Virtual Keith to the better, faster software sought by Harry Mott, it may be the attempt to render and re-imagine, in greater detail and from any perspective, the very things around us. That means graphing stock market data as pulsating cones or, as students from MIT demonstrate in the mesmerizing Alpha Wolf, creating a biologically sound simulation of a wolf’s lifecycle among the pack, to be played as a video game.”

**LA Times, 8/17/01, “Where the Mona Lisa Meets the Motherboard” [Gaslin 2001]**



## 6 Discussion

While creating, exhibiting and evaluating AlphaWolf, a number of interesting phenomena have come to light. First, the process has provided insight into the creation of directable social characters. Second, the system has offered an example of how to craft installations around these characters, addressing a number of thorny issues in interface design, game dynamics and related issues. Finally, AlphaWolf has proven to be effective at enabling the three kinds of social relationships mentioned in the Introduction – machine-machine, human-machine, and human-human mediated by machines. This section discusses each of these areas in an effort to work out what was “right” about AlphaWolf.

### 6.1 *Directable Social Characters*

One of the most central issues in creating the system described in this thesis was that of making social characters who could be directed by a human participant. Various aspects of social and directable characters are discussed below.

#### 6.1.1 Necessary Components

The success of AlphaWolf at enabling the formation of the relationships among computational entities and people confirm that the AlphaWolf system includes a sufficient set of representations to enable the creation of the kinds of relationships that I mentioned in the introduction. Nevertheless, it is hard to tease out the necessity of every element of a system as large as the Synthetic Characters’ code base (e.g., the relationships could still have been formed without the charcoal renderer, but the installation would have been an inferior product without it). Therefore, efforts were undertaken only to confirm that the key elements of the CSEM mechanism – emotion, perception and learning – were necessary for the social relationships to be created. The various failures of the three alternate mechanisms – Fixed, Random, and Emotional – confirm the necessity of each of these three key parts.

The failures of the Emotional and Random algorithms demonstrate that perception is necessary. Without perception, wolves would be forced either to rely on their current emotional state to inform their behavior (if they had an emotional state), or to pick their interaction style at random (if they did not). As the synthetic experiments show, there is no significant persistence of relationships with either of these algorithms.

The failures of the Fixed algorithm demonstrates that a dynamic emotional state is necessary. Without the ability to respond emotionally to their interactions, the wolves are fixed in their interaction roles. In the Button/Fixed runs of the human user studies, it was impossible for subjects to encode relationships reliably. In the resource exploitation simulations, the Fixed algorithm was not dynamically reconfigurable (although it did show clearer relationships than the AlphaWolf algorithm). Therefore, although the Fixed algorithm could be appropriate for certain kinds of problems, its shortcomings demonstrate that a dynamic emotional state is a necessary component of the CSEM mechanism.


The third necessary part of the system is learning. All three alternate algorithms lack learning; without it, the virtual wolves do not exhibit relationships that persist over long-term series of interaction. Although long-term social relationships may not be necessary for certain kinds of natural social phenomena (as described in the Related Work chapter), the lack of learning eliminated the CSEM’s ability to transmit social relationships and caused the pack of wolves to be unable to form a directable distributed dominance hierarchy.

To summarize, the AlphaWolf system as a whole is sufficient to create synthetic social relationships, and emotion, perception and learning are all necessary parts of that system.

### 6.1.2 Clear Emotions

In their interactions with the AlphaWolf system in its various forms, people tended to like wolves who showed strong, clear emotions. The Emotional algorithm, which performed best of the four social relationship mechanisms, allowed users a lot of control; after they dominated or submitted to another wolf, the effect of that interaction was clear in the behavior of both pups. In addition, it caused wolves to spend a lot of time performing the extreme versions of their animations. These extremes were the animations that had been hand-animated by our professional animator, and were therefore the strongest parts of the dynamic range. The blended animations in the middle of the range were by necessity a bit muddier, being a combination of several hand-made examples. The lesson learned here is that strong, clear, responsive emotions help make good characters and installations.

This lesson is one that Disney animators know well. To make compelling characters, have strong emotions (see Figure 6-1). Building on a history of strong (if simple) characters in theater and other previous entertainment media, early animators embraced the notion of a character who shows only one emotion at a time.



**Figure 6-1: The gray pup walking sadly.**

Nevertheless, human actors in movies are usually lauded for more complex performances. If a human actor gave a performance like that of a Disney character, he would be called “flat” or “histrionic” or “melodramatic.” More subtle and complex interactive computational characters will hopefully join the ranks of human actors in their complexity, rather than continuing in the simplistic vein of traditional animations. In order to have “staying power” during the long time-scales in which interactive characters will be experienced (e.g., 40 hours for a video game, hundreds of hours for a “personal digital assistant”), those characters will need to be capable of subtlety. However, this subtlety should be drawn from a wide, clear, expressive, dynamic range.

In fact, if there exists a wide range of expressive behavior, it is easy to make a shorter time-scale experience with it. By causing the entity to move more rapidly through its emotional space, and stay near the hand-animated extremes of its repertoire, a designer could cause the character to exhibit a more spiky emotional profile. Regardless of the length of the interaction, an installation should explore its characters' expressive ranges; a longer interaction has more time to explore those ranges in a leisurely fashion.

### 6.1.3 Controllable 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.

As I described in the Implementation Chapter, 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. Each Action comes bundled with a number of TriggerContexts, which determine when the Action happens, and several DoUntilContexts, which determine when the Action finishes. An Action may also have an object, which is the target of the action. For example, a sleep action's TriggerContext might be “when fatigue is above a threshold”, its DoUntilContext 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 DoUntilContext 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 TriggerContext 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.

### 6.1.4 Directability

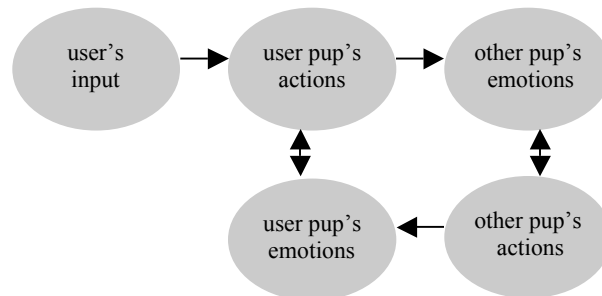
There is some essential interplay between action and emotion in semi-autonomous characters. As Antonio Damasio points out, emotion is central to the way people (and animals) decide what to do [Damasio 1994]. How can we make a user responsible for the actions of a character, and yet allow emotion to play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

Our way of causing a user to behave in a manner consistent with the desires of his virtual wolf is to streamline the interaction – make it 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 fact, if the user finds his pup walking slowly, he can choose to change his direction to the pup; if the new direction matches the pups desires, it will switch to the faster gait.

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; this section of the behavior system also enables a pup to determine autonomously whether the user’s suggestion matches what the pup 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.

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 more dominant individuals. Therefore, the pup may look away, occasionally glancing back at the dominant individual.



**Figure 6-2: The flow of control among a user and two pups.**

Figure 6-2 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 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 its submissive relationship to the other pup, does so in a crouched posture, looking away as it 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. This example demonstrates that, through the AlphaWolf system, a participant can direct a virtual wolf to do something that is essentially "against its will", and yet the character will do it and still remain in character.

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 and less dependent on the user as the interaction proceeds.

In creating a directable character, it is relevant to consider what elements of it should be autonomous, and which should be controlled. In AlphaWolf, the actions were controlled, and the emotions autonomous (albeit influenced by the actions). The interface should help make the distinction clear; a natural interface (e.g., gesture recognition, microphone) might be more appropriate to high level control than an artificial interface (e.g., a keyboard). In my experience, participants don't mind having limited control as long as they know that they have only limited control; problems arise when participants expect to have complete control, and then feel stripped of it when the limitations are revealed.

## ***6.2 Character-Based Interactive Installations***

Once we have directable social characters, the battle is only half won. Situating these characters in an installation that shows them off, focuses participants' attention on them, and otherwise supports the characters, is another great challenge.

### **6.2.1 Expectation Management**

People bring a lifetime of conceptual baggage to a new experience such as a social character-based interactive installation. In the last year, a person may have watched her child playing video games, seen the movie *A.I.*, read National Geographic Magazine, been to an art museum, and flipped past any manner of animated or real wolves on television. While it is not possible to make an interactive installation that takes into account every possible allusion that every participant may draw, it is essential that the installation have some awareness of the cultural context in which it is being presented.

Taking computer games as an example, the game *Doom* features a first person camera angle and a gun centered at the bottom of the screen. If AlphaWolf had featured a first person camera and "your" wolf paws walking at the bottom of the screen, it might have predisposed participants to thinking of the installation in terms of that game. Each of the choices made in AlphaWolf, from

the rendering style to the interface to the autonomous cinematography to the sound effects, was crafted to make certain allusions and not others. Ultimately, we wanted to cause people not to think of the installation as the newest “Wolf Hunter” game, where the goal is total domination of the pack. Rather, we wanted them to think of it as an investigation of social behavior, relationships, and group dynamics.

The notion of expectation management also factors prominently into the question “What makes a good interactive experience?” Ultimately, no one thing makes a good interactive experience. What makes a good interactive experience is the satisfaction of certain expectations and the creation of others. For example, when a person sees a computer mouse near a screen, they expect that it will have some effect. The mouse in front of the AlphaWolf screen has a significant effect, in that it causes the wolf to look at and run to the place where the person clicked. The wolf does so immediately, and people seemed satisfied with that control.

On the other hand, people do not habitually expect a cartoon wolf to howl back when they howl at it. The first time people howled into the mike and had their wolf howl back, it often changed the way in which they expected to interact with the computer.

Creating a good interactive installation, then, is an art. Just as a painter uses form, composition and color to lead a viewer’s eye around a canvas, an installation designer crafts elements to lead a participant’s understanding of the experience. A good installation is not necessarily one in which people have a lot of control, or one in which they feel immersed, or one in which they identify with the characters, but it *can* be any of these things. Making the participants feel something they’ve never felt before, or see the world in a new way, is the goal. In order to do that, the designer must have some idea of what participants have felt before, and how they see the world.

## 6.2.2 Game Play

Interactivity and interface are parts of a more nebulous topic – game play. Why are some games better than others? Why do some hold our attention for hours (even if very simple or repetitive), and other get boring? Looking at an interactive installation as a game offers up some ways of thinking about its design.

### **What is the genre of the game?**

There are many different genres of games, from racing games to puzzle games to hunting games. For AlphaWolf, the genre into which it most closely fits is that of social games – games like *The Sims*, *Creatures* or *Dogz*.

### **What is the structure of the game?**

Some games go on and on (e.g., *Sim City*), others have a number of discrete levels (e.g., dungeon-exploration games), others are short but addictive (*Minesweeper*). For AlphaWolf, we wanted a short interaction that nevertheless felt like part of a longer whole. The life cycle of an animal, with each generation being a single interaction episode, fit this structure.

### **What is the role of the player?**

The player of a game can take on a number of roles, depending on the style of the game. The player sometimes acts as omniscient overseer, other times directs several characters, other times plays a single character. The role of the player was tricky for AlphaWolf, in that we wanted to have a one-to-one correlation between participants and their characters, but still have those characters exhibit some independence. Ultimately, we settled on having the player be a “guide” for a semi-autonomous wolf pup.

**What are his or her controls?**

The role of the player is intimately linked with the controls that he or she has. The balance that we wanted to strike between directability and autonomy made some kind of high-level control seem most appropriate; the focus on social behavior suggested an interface that evoked social phenomena. The microphone seemed ideal as the primary mechanism for control.

**How does he or she win?**

People come into a game with an expectation that they can somehow win, or at least finish, the game. Since the AlphaWolf installation was designed to have multiple players interacting at once, we did not want all but one to be losers at the end. Therefore, we tried to frame the piece as “finding your pup’s place in the social order of the pack”, rather than “achieving dominance.” With that common agenda for all three interactors, the challenge was not to defeat the opponents, but to collaborate with them and work out the hierarchy.

Each of these choices predisposed people to view AlphaWolf for what it was – an interactive exploration of social relationships. Using the social development of a litter of pups as the “script” for the unfolding of events gave people an intuitive grasp of how the interaction was meant to proceed. Because they knew what to expect, at least to a certain extent, they were able to like the game. It tried not to promise anything it couldn’t deliver.

### **6.2.3 Interactivity**

Interactivity is something that people are becoming more familiar with in the world of computing, in particular through games and the World Wide Web. The human user study that I conducted examines some of the specific effects of interactivity on participants’ subjective experience of the installation.

The results from the study show that interactive versions of AlphaWolf offered subjects an increased feeling of control over the wolves’ behavior and a more immersing experience. However, interactivity also reduced the clarity of the topic presented, and made the wolves less realistic. In addition, interactivity was less significant than cinematography for building a relationship between a user and the leading pup.

Some of these effects may reflect the novelty and difficulty of having to learn to play a game, however simple, in four minutes. The fact that novelty was sufficient to swamp the effect of the different social relationship algorithms suggests that it could certainly be significant enough to sway people’s subjective experiences in these topics as well. It is interesting to consider that the same novelty effect almost certainly occurred in the early days of cinema. Gamers who participated in the AlphaWolf study understood it almost immediately; how many years will it be before interactivity is as pervasive as moving pictures are now?

Considering that novelty was a great part of the reason that people preferred the second run over the first, I imagine that clarity of relationships would have increased as users continued to interact and the interaction paradigm became intuitive. However, in a setting where the first few minutes are all that matter (for example, at a museum installation) novelty will always have a significant impact, as nearly all interactors are novices. It’s important not to underestimate the effect of novelty, or to forget about it because an installation isn’t novel at all to its creators. (Of course, interactive experiences will lose their overall novelty as they become more ubiquitous in the world’s media.)

Another lesson regarding interactivity is that, depending on the topic at hand, an interactive experience is not necessarily better than a non-interactive experience. Since interactivity reduces the similarity of the virtual wolves to real wolves, a short-interaction virtual wolf pack might be a great game, but do little to teach people about wolf social behavior. Again, though, this effect might be diminished in a longer piece, as immersion increases. In addition, an interactive installation, even a short one, might be a way to get people interested in a topic, so that they'll stay around long enough to learn about it from other media. I'll discuss these topics further in the Applications section of this document.

The impact of the two different interactive paradigms – Button and Microphone – was significant as well. Perhaps the Microphone paradigm, while not appropriate for the shy in an audience and invariably taking a bit longer to master, provided richer rewards in the long-term.

#### **6.2.4 Social Is Intuitive**

The ease with which people from diverse cultures, backgrounds and age groups figured out the AlphaWolf installation supports the notion that people have an intuitive grasp of social phenomena. In both the AlphaWolf installation at SIGGRAPH and the human user studies, people rapidly grasped the interaction paradigm. Social interactions appear to make sense to people. People have an intuitive grasp of dominance and submission. Nevertheless, I come back to this topic in section 6.2.7 below.

People's intuitive ability suggests that social structures, and in particular dominance structures, might be useful constructs for helping people understand the workings of large, complex systems (e.g., multi-robot systems). A hierarchy of command, as already exists in many corporate and military organizations, could extend into the robotic system, or be paralleled by it. Many different interactions among groups of entities could be cast as dominance interactions, as a means of understanding them in human-accessible terms.

#### **6.2.5 Social Is Effective**

In addition to creating a compelling experience for human interactors, the AlphaWolf mechanism is good for making controllable multi-agent systems. The distributed social arrangements depicted in the resource exploitation simulations show that AlphaWolf makes dominance hierarchies that can be readily re-ordered. In addition, this re-ordering can be controlled by affecting the behavior of a single individual. As a multi-agent or multi-robot system got bigger, this point would become more and more important. While a fixed hierarchy might provide slightly better performance on a limited set of constrained tasks, AlphaWolf hierarchies would be better on open-ended tasks, tasks that need to be controlled explicitly, and tasks in which entities learn about their world.

Some of the most interesting applications of the AlphaWolf mechanism for multi-agent systems and multi-robot systems come when social structures are attached to real-world activities like recharging. When dominance affects the way in which entities interact with other parts of their environments, it becomes a construct that is useful for more than just enhancing human understanding of the system. A future extension of this work would be to see if dominance hierarchies could increase the refueling efficiency of a group as a whole, with all individuals refueling as rapidly as possible.



### 6.2.6 The Implications of Dominance

The AlphaWolf installation focused primarily on dominance behaviors; this focus has implications for how people perceived the work, and also for its extensibility to other kinds of social relationships.

With regard to people's subjective experience, dominance may have played a role in the gendering of the wolf pups. Dominance, at least in wolves, is perceived to be a more masculine domain (e.g., the Alpha male), so it is unsurprising that people projected a masculine gender on the genderless wolf pups, even though real wolf pups of both sexes engage in the kind of dominance interactions depicted in AlphaWolf. It would be interesting to discover what behavior patterns would cause people to view the wolf pups as female.

Another implication of the dominance focus of AlphaWolf is that it plays into some of the deepest fears of technology that people harbor. Machines traditionally do not compete with people for the same resources. (There are certainly some exceptions, as a factory worker who has lost his job to automation will admit.) Machines, in fact, have not tended to be Darwinian in their existence as a group; they do not autonomously compete with each other for reproductive opportunities, etc. One of the great latent fears of technology is that machines will compete with us, and in fact *out*compete us. The futuristic world that spawned the assassin robot in *The Terminator* was one in which machines had deemed it appropriate to eliminate humanity. The AlphaWolves, on the other hand, attempt to introduce people to the idea of computational systems as social companions that are not all-powerful and malevolent. While interactions between machines and people do not need to revolve around dominance, dominance is a domain with which people have a lot of experience; therefore dominance is a domain in which people might come to accept other social entities even though they are, in the case of AlphaWolf, computational.

Ultimately, dominance and submission are just one type of social behavior. Dominance is a good model for certain kinds of interactions between people and machines, but certainly not for others. While a dominance hierarchy might be a clear way of understanding the interactions among a group of autonomous robots, it is probable that people do not want to be challenged by their household machines for dominance. The goal of this project has been to show that social relationships can exist among computational entities, rather than to propose that machines with a penchant for dominance are the next big thing in household appliances. The larger issues of interactivity and its interrelation with autonomy are important regardless of the particular kind of social relationship.

### 6.2.7 Implicit Biases

While the research goal of this thesis has been to make an explicit computational model of social relationships, a number of implicit biases have been built into the system. These biases lie at the heart of the game dynamics that lead people into certain interactions and away from others.

The most central bias is that social relationships are the most compelling aspect of the virtual world. People who seek out items and interactions beyond the social sphere will be sorely disappointed. This bias rapidly became explicit if a participant spoke with one of the creators, but in the installation itself the conspicuous absence of anything but the wolves and a few ethereal trees led people to attend primarily to the social phenomena.

Within the social interactions, if two people got in a growling match, loud, gritty voices were favored over higher, weaker ones. However, there was a hint that we sometimes gave to people

who seemed to be getting the short end of the interaction; blowing into the microphone caused it to clip, eliciting a vigorous growl from that pup.

Growling was actually somewhat discouraged by the installation. The growl noise is quite hard on the throat, and can be interrupted readily by talking, laughing, or any non-growl noise. It took quite some persistence to growl steadily for the full five minute interaction.

In terms of having a strong say in the spirit of a relationship, initiators were favored. Because a wolf's actions had an effect on both itself and the other wolf, individuals who started interactions tended to have an easier time of dictating the direction that the relationship would take than those who waited to be approached.

A more subtle bias that we made sure to support is that just dominating is not necessarily the best strategy for achieving status. While a "growl all the time" strategy will certainly result in a more dominant wolf than a "whine all the time" strategy, a mixed strategy of growling and playing seemed to be the best recipe for overall dominance. It was important to us that AlphaWolf not degenerate into a war game for leadership of the pack. The adults, who trounced any pup who growled at them, helped keep certain growl-happy participants in line. (In fact, a wise dominance-seeker might wait until an adult disciplined a growl-happy pup, and then coat-tail off the trouncing. Running over and giving a little growl at the now-submissive pup would establish dominance without a prolonged growl-off.)

While we tried to create an installation that cut across international boundaries (e.g., no text or spoken language), there may be certain implicit biases in the system toward American or European ways of thinking (since the production team was largely from those cultures). Masuda and Nisbett [Masuda 2001] point out that Americans understand scenes in a different fashion than the Japanese, with Japanese taking a more holistic view, and Americans focusing on animate entities. I did not conduct any studies to ascertain the presence or absence of any cultural bias in the AlphaWolf project. It would be interesting to determine if the attributes of the AlphaWolf social relationships resonate more fully with one culture than another.

### 6.2.8 Simplicity

The AlphaWolf representation of social relationships is very simple, and the mechanism by which it is applied is also. Nevertheless, while it may not capture all the complexities of wolf social relationships, and falls well short of human relationships, it is still clearly *recognizable* as a social relationship by people.

The simplicity of the model has several ramifications. First, it suggests that it could be applied to a wide variety of domains. Because it does not presume too much about the entities that will be interacting or the ways in which they will interact, the model should be readily able to be modified and extended. The model's simplicity also suggests that it should be fairly robust. While it does not cover the entire range of social phenomena, it should capture those that it *does* cover in a adaptable and extensible way so that it might be extended in the future. Finally, the simplicity of the model hopefully means that it will scale well, and can be applied to far more complex systems.

In terms of installation design, the simplicity of the AlphaWolf installation and wolves intentionally leaves room for participants' imaginations. By underspecifying certain aspects of the world, through the sketch-like rendering style, the simple behavior model, the barren space, we hoped to draw people in to the world, rather than assaulting them with too much detail.

## 6.2.9 Real Time

The fact that AlphaWolf is a real time system is very important to its success. Previous research projects have explored ideas of multi-agent coordination by means of social memory in simulation (e.g., [Hogeweg 1988; Hemelrijk 1996]). These implementations, though, lack human interactivity, which is one of the central themes in this dissertation. Making systems that can interact with humans socially, and can be directed by them within a familiar social paradigm, is a significant step beyond these previous projects. The real time nature of our system makes this human interaction possible.

## 6.2.10 Stories

As I mentioned previously, stories often focus on social and emotional phenomena. A group of social synthetic characters such as the wolves, therefore, can be seen as an easy breeding-ground for stories. People who interact with AlphaWolf sometimes like to tell these stories, to each other, to us, or to the questionnaires in the user tests when they asked for “any additional thoughts.”

A story is essentially a scenario populated by characters who undergo emotional change. By creating virtual wolves who are able to undergo emotional change and putting them in a common virtual world, we are able to establish situations in which stories emerge from the ensuing interactions. Since our characters are still very simple (compared to real people or animals), the stories that emerge are also very simple. However, as virtual characters become more complex, we hope that the stories that emerge from their interactions will become more interesting.

One of the features that separates an interactive story from a “regular old story” is the ability of one or more participants to influence the course of that story. In AlphaWolf, that is exactly what happens. Participants assume the role of one or more of the virtual characters, and can thereby interject their authorial control over the situation. By situating participants in starring roles (at least to their own automatic cinematographer), AlphaWolf creates interactive stories in which those participants are *central* to the events and emotions that occur.

One of the ways in which we set up the story of AlphaWolf is by giving it a beginning. Rather than starting out *in medias res*, the installation starts off with the wolves asleep, needing to be woken up by participants in order to engage in interactions with their pack mates. The waking up helps engage the participant and forge a connection between participant and pup.

While we do have a simple ending in place at the end of the five minute interaction, where all the wolves rally around the alpha pair, there might be room for improvement in closing the interaction. Just as AlphaWolf opens with a personal connection between each participant and his or her wolf, it might be a strong closing to have each pup engage its user, making eye contact through the screen, and behaving in a manner that reflects the personality that has developed in it over the preceding minutes.

## 6.2.11 Cinematography

One of the more surprising results that appeared in the user studies was the profound impact that cinematography has on the relationship between a human participant and the character that he or she is playing. Interactive cinematography has a clear effect on the way people perceive virtual characters. As the data from the user tests show, participants like the pup that the camera focuses on more, and feel a stronger identification with it. Interactivity (i.e., the ability to control that character) by comparison, had much less effect on either liking or identification.

Why does cinematography have such a significant effect? It has an effect for the same reasons it has an effect in film and television. People are used to having their attention directed, and readily take cues from the camera. The camera usually watches the actions of the protagonist; if it's looking at the nemesis, the cinematic treatment of that character generally makes that perfectly clear. The interactive cinematography system in AlphaWolf was built to allow people to see the social relationships involving the gray wolf, and, more importantly, to see them from the point of view of the gray wolf. We're looking over the shoulder of the gray wolf (see Figure 6-3); we're "on its side."



**Figure 6-3: People like and identify with their gray pup.**

The lesson from these results are to make sure to pay attention to interactive camera work. In this case, it worked strongly in favor of the gray pup. Bad cinematography, on the other hand, could presumably do as much harm as this camera did good.

Perhaps a more interesting question, though, is why interactivity does *not* have more of an effect on liking and identification, and was rendered almost unnoticeable by the effect of the camera, when interactivity had such a strong effect on other aspects of users' subjective experiences. The two possibilities are that interactivity simply has little effect, or that there is some positive effect but an equivalent negative effect that cancels it out. In the latter case, the sense of identification and liking are increased by the act of directing the pup and therefore building a rapport with it, but that it also takes away from the excitement of finding out what's going to happen next – the gray pup might seem less interesting because it is directable. Regardless of why interactivity does not have more effect, it is clear that cinematography has a profound impact on an audience's impression of a character, interactive or not.

### **6.2.12 Sound**

Although sound was not one of the specific areas studied in the user tests, sound effects played a major role in those tests and in the AlphaWolf installation at SIGGRAPH. The several times we have run the wolves without sound, people don't become engrossed. It doesn't feel "finished." While a few people complained about the sounds, I believe that they are crucial to the installation's success.

As anecdotal evidence of the importance of sound, the moment that people have their first realization of their control over their wolf is when they howl into the microphone and the pup howls too. This moment, more than any other, provoked looks of wonder. It is my belief that a wolf who simply played a howl animation, with no howl sound, would not have “grabbed” people nearly as well.

Sound and cinematography, which have been two of the main “supporting technologies” in the Synthetic Characters’ installation-building tool kit, are both crucial to the success of an installation.

### **6.2.13 Physical and Contextual Setting**

While the vast majority of topics covered by this thesis involve things going on “inside the box”, the physical setup of an installation does make quite a bit of difference, and bears mentioning again. As I described earlier in the chapter about the AlphaWolf installation, our change of the physical arrangement of the installation appears to have changed the way in which people interact with it and with each other.

In addition, contextual information has a significant impact as well. In the user tests, starting people off with the video of real wolves predisposed them to think about the virtual wolves *in terms of* real wolves, rather than just as a game. By framing the virtual wolves appropriately, people thought to look for correlations that they otherwise might have missed.

The contextual framing of an installation is a much broader topic than can be covered here. As I discussed above, the conceptual baggage that people bring to an installation (e.g., is this a research experiment or an interesting diversion at a graphics conference) and the set of expectations that are created in a user by the elements of the installation itself (e.g., what allusions and references are included in the work) both color how people will perceive every aspect of the installation. This topic is discussed in greater depth in a previous paper by several members of the Synthetic Characters Group [Blumberg 2001].

One interesting distinction lies between the AlphaWolf installation at SIGGRAPH and the version employed in the human user study. At SIGGRAPH, the microphone interface, the large audiences and the networked pack worked together to create a vigorous social setting for the installation, with people excitedly engaging the wolves and each other. In the human user study, though, where each person was seated in a room alone throughout their interaction, the social context could not have been nearly as central to participants’ subjective experience. I expect that surveying people in the context of SIGGRAPH would have revealed much stronger support for the microphone interface than the isolated studies did.

### **6.2.14 Process**

The AlphaWolf project is changing the way we build characters – rather than assembling finished adults, we are now creating young pups with certain instinctive behaviors and the ability to grow up and integrate into the social context of their pack. This gradual shift in the process by which we work is indicative of our emphasis on learning and real-time adaptation as crucial parts of a synthetic character that is to be seen as intelligent or life-like.

This change has caused a corresponding increase in our emphasis on visualizers for the various elements of our system. If a system is going to change its values or structure during the course of its execution, it is necessary for the developers to be able to find out how it’s changing and why.

The implementation of this document has included various samples of the visualizers that we use in our system.

Effective visualization tools are an essential component of building increasingly complex systems. While it is possible to keep all aspects of a simple character, or even of an entire simple world, in one's head at the same time, AlphaWolf rapidly moved beyond that level of complexity. Future projects, which will inevitably be more complex in some way, will continue to rely heavily on visualizers during the development process.

In addition, good visualizers can be used to help explain to interested people how and why the characters' decisions are being made. At SIGGRAPH, for example, we had an extra machine that was just showing the visualizers for the other characters. People who were interested in what was "under the hood" could examine the structures that we used to craft the behaviors and relationships among the wolves.

### ***6.3 Three Kinds of Relationships***

In the Introduction, I mentioned three kinds of relationships that are relevant to this work – human-machine, machine-machine, and human-human. In this section, I demonstrate that each of these three kinds of relationships were formed through the AlphaWolf system. In each of these cases, I refer back to the definition proposed in the introduction: "a learned and remembered construct by which an entity keeps track of its interaction history with another entity, and allows that history to affect its current and future interactions with that entity." Breaking down this definition, we must satisfy the following five parts of it:

- "learned and remembered construct" – there must be persistent storage of information about social partners, and that information must be acquired over the course of one or more interactions
- "entity" – there must be a notion of an entity
- "interaction history" – interactions must have an effect on an individual, and in turn on the remembered construct
- "another entity" – an entity must be aware of its social partners as entities that persist between interactions
- "affect its interactions" – the remembered construct must have an impact on behavior

For each of the three types of relationship, I confirm that each of these conditions was satisfied. As a caveat before starting this analysis, I will admit that definitions are elusive and sometimes misleading. While I try to present evidence that relationships were formed, and that many of these relationships involved people, it is inappropriate to imagine that the relationships described here approach the complexity and interest of a full human (or even lupine) relationship. These analyses merely confirm that the interaction patterns seen in AlphaWolf appear to reflect the above definition.

#### **6.3.1 Relationships between Machines and Machines**

In all three implementations – the SIGGRAPH installation, the human user studies, and the foraging simulations – the virtual wolves demonstrably formed relationships with each other. Table 6-1 demonstrates that each wolf satisfied the five parts of the definition.

While the relationships between pups are admittedly very simple, they nevertheless satisfy the criteria in the definition. In the Future Work chapter, I describe a number of ways in which the virtual wolves could be made to form more complex and interesting social relationships. The fact

that the virtual wolves formed relationships with each other will be important in the next two section, where it will form the basis for the analyses of the other two kinds of relationships.

Part of Definition	Validation
“learned and remembered construct”	The wolves started off identical, and their relationships were learned over the course of the interaction. The CSEM representation of one wolf could then be loaded into another wolf in a different run and decoded accurately by novice human participants. Therefore, the CSEM representation served as a remembered construct that stored the wolves’ relationships.
“entity”	Wolves and people formed different relationships with the black pup and the white pup. The autonomous wolves and those directed by another user exhibit behavior that is unpredictable within certain constraints. Therefore, wolves can be seen as distinct entities.
“interaction history”	The interactions that wolves engaged in affected their emotional state, and thereby influenced the internal representations of their social relationships.
“another entity”	Wolves could distinguish between their siblings and form different relationships with them.
“affect its interactions”	The CSEMs affected the wolves emotional states, and thereby affected their choice of behaviors (in fully autonomous wolves) and the style in which they performed those behaviors (in all wolves).

**Table 6-1: The virtual wolves satisfy the definition of entities forming social relationships.**

The relationships among wolves tend to be reciprocal; if wolf **A** believes itself to be dominant to wolf **B**, then **B** probably believes itself to be submissive to **A**. While it is possible to have relationships that momentarily violate this tendency, the positive feedback system of CSEMs causes relationships of equal standing to be uncommon and transient, as slight variations become amplified. Other phenomena such as cyclical sets of relationships ( $A > B > C > A$ ) are possible, but tend to be “ironed out” into linear arrangements by the system.

### **6.3.2 Relationships between Humans and Machines**

During the interactive runs of the user tests, people took on the role of the gray pup and formed relationships with two computational systems in the shape of the other two pups. This description points out that there were in fact two different kinds of human-machine relationships formed. The first kind was the relationship between the human and the gray pup that he or she was controlling. The second kind was the relationship that the participants formed with the other wolves when they were playing the role of the gray pup. Each of these will now be addressed in turn.

#### ***6.3.2.1 Relationship between Participant and Gray Pup***

Each participant in the user tests played the role of the gray pup. Table 6-2 demonstrates that both the person and the gray pup formed relationships with the other over the course of the interaction. However, whether or not these relationships qualify as social relationships or merely patterns of repeated interactions is unclear.

Part of Definition	Validation – Human	Validation – Gray Pup
“learned and remembered construct”	People started with no knowledge of the various pups, and learned about them during the interaction. People were able to answer questions about the gray pup after each run had finished. Therefore, they had formed some kind of remembered construct of that pup.	The relationships formed by the gray pup with its pack mates through the CSEM mechanism could also be seen as relationships with the participant. An example of this is that pups changed the way in which they obeyed the participant (e.g., walk vs. run) depending on the kind of interaction the pup had come to expect from the user.
“entity”	Humans are autonomous entities.	Many people treated the gray pup as a separate entity from themselves (e.g., by using the third-person pronoun when referring to that pup). Some, though, did not, subsuming the pup into themselves.
“interaction history”	People had different opinions of the gray pup in different runs. Therefore, they had distinct interaction histories with each version of the gray pup that they directed.	The gray pup’s “personality”, i.e., its overall dominance profile, shows that it stored some element of its interaction history with the participant.
“another entity”	Many people treated the gray pup as a separate entity from themselves (e.g., by using the third-person pronoun when referring to that pup). Some, though, did not, subsuming the pup into themselves.	The action/emotion split, through which a participant controlled the actions of the wolf and the wolf itself controlled its emotional state and relationships, makes a clear distinction between the wolf as an entity and the participant as a separate entity.
“affect its interactions”	People’s memories of their interactions with the gray pup influenced not only their behavior during each run, but also their responses on the questionnaires.	The CSEMs that the gray pup had formed on account of the participant’s directed actions affected its behavioral choices and the style in which it performed those behaviors.

**Table 6-2: Humans and the gray pups (barely) satisfy the definition of an entities with social relationships with each other.**

Despite the above validations, there were significant weaknesses in this analysis. The first problem lies in the lack of any explicit representation of the participant on the part of the gray pup. While the collection of all the CSEMs does serve to model the net interaction history taken by the participant, these actions were taken towards the other wolves in the pack, not towards the gray pup itself. This problem calls into question whether the pup formed a relationship with the participant. Also, the fact that many people used the first person pronoun with respect to the gray



pup called into question the existence of the gray pup as a separate entity. This problem calls into question whether the participant had formed a relationship with the gray pup. In fact, there is probably a gradient from social relationships to other sets of repeated interactions, and this “relationship” falls somewhere in the middle. As we will see in the next section, the support for the relationship between the participant and the virtual wolves whom he or she was not directing is much stronger.

### **6.3.2.2 Relationship between Participant and Other Wolves**

Participants in the user studies clearly formed social relationships with the other pups that they interacted with through the gray pup, as is shown in Table 6-3. As Table 6-1 above showed, those pups formed relationships with each other and with the gray pup. Since the participant was in control of the actions of the gray pup, those relationships were in fact with that person (although only with that person in the guise of the gray pup).

Part of Definition	Human	Machine
“learned and remembered construct”	People were able to answer questions accurately about the different pups after each run had finished. This demonstrated that they had formed memories of them.	The fact that the pups’ relationships could be loaded into another run and decoded accurately by people demonstrates that they had formed memories of each other.
“entity”	Humans are clearly individual entities.	People showed no hesitation at treating pups as entities (e.g., answering questions about them).
“interaction history”	People behaved differently toward the two different pups, and showed clearly different opinions about them. Their interactions with the two pups left them with different (and often strong) viewpoints towards each.	In the AlphaWolf runs, the interactions that pups engaged in affected their emotional state, and thereby influenced the internal representations of their social relationships.
“another entity”	People’s ability to form opposite relationships with the two pups clearly shows that they could tell them apart and deal with them as separate entities.	Pups could distinguish between their siblings and form different relationships with them.
“affect its interactions”	People’s memories of their interactions with the pup influenced not only their behavior during each run, but also their responses on the questionnaires.	In the AlphaWolf runs, pups changed their behavior based on their interactions.

**Table 6-3: Participants formed relationships with the wolves whom they were not directing.**

People’s success at encoding relationships, even in the first run, shows that they already knew how to form these relationships. Subjects’ success at decoding shows that the other half of the equation is also in place – people could understand the wolves’ expressiveness.

While the user tests demonstrate these relationships with a certain degree of quantitative validity, the participants in the AlphaWolf installation confirmed these trends on a much larger scale.

Since the carcass-exploitation simulations do not include people, it might seem that they do not have much to add to the notion of relationships between humans and computational entities. However, since there is the element of directability in those simulations, they do offer some insight into how humans might participate (albeit indirectly) in the relationships between machines and other machines.

### **6.3.3 Relationships between Humans and Humans**

A bit less clear is the case for AlphaWolf helping to build relationships between pairs of people interacting with it. However, if two virtual wolves can be said to have formed a relationship (machine-machine), and it is possible to extend this analysis to include a person playing the role of one of the virtual pups (human-machine), then it seems reasonable that both wolves could be “swapped out”, and the two people directing them could be said to have formed a social relationship. Since there were no people in the foraging simulations and only one person at a time in the user studies, the AlphaWolf installation is the only one of the three where this claim might be made strongly.

At the other end of the interaction, people’s encounters as synthetic wolves often influence their ongoing relationships in real life. People strike up conversations after runs of the wolves (the wolves “break the ice”), having common ground to discuss. While I have no evidence that the effect of the wolves on this kind of relationship is much stronger than any other co-operating experience (e.g., playing pool, riding a ski lift), it does seem that the topic of social interaction is fresh in people’s minds. Perhaps this is why people go to see movies on first dates – movies are usually about social and emotional occurrences, and can therefore act as a good ice-breaker for topics appropriate to human relationships. Whereas I do not believe that most video games are as appropriate a pastime as a movie for a first date, I would hope that an interactive experience like the virtual wolf pack might be even better than a movie. I’ll discuss this further in the Applications chapter below.

If people did form relationships with each other through the AlphaWolf installation, they were low-bandwidth relationships. As I show in the next chapter, one of the significant limitations of this AlphaWolf mechanism is its inability to represent the complexities of human social interaction.

## **7 Future Work and Applications**

Minsky defined artificial intelligence as “the science of making machines do things that would require intelligence if done by men.” [Minsky 1968] Similarly, we might define social synthetic characters as the science of making machines do things that would be called social if done by people or animals. Given this definition, the virtual wolves in the AlphaWolf installation are just one example of social synthetic characters, and a simple one at that. In this chapter I discuss the limitations of the system, ways that the model presented here could be used in larger-scale and more complex projects, and a selection of potential applications for the group of ideas presented here.

### ***7.1 Limitations***

Despite the relative success of the AlphaWolf system, there are a number of significant limitations to the system and the relationships that it forms. Many of the shortcomings mentioned here will be addressed in a variety of forms in the Future Work section as well.

#### **7.1.1 Emotional Simplicity**

The first and clearest limitation is in the area of emotion. While the single axis of dominance has served admirably as a simple, first-pass emotion system, we have always intended to have the system be extensible to more complex emotional models.

This limitation on the emotional complexity has several ramifications. It greatly reduces the available bandwidth that can be transmitted by the relationship mechanism (see section 7.1.4). It prevents the system from capturing certain elements of wild wolf social behavior (arousal, for example, is central to the process by which adult wolves feed pups [Mech 1988]). Finally, it does not fully explore the potential connections among emotion and social behavior.

#### **7.1.2 Perceptual Simplicity**

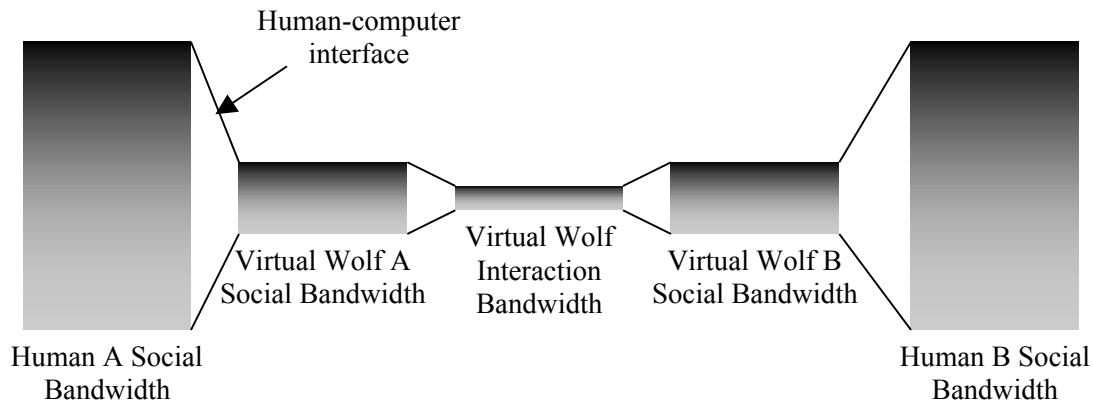
A second significant problem lies in the area of perception. One of the phenomena that we would hope that this mechanism could explore is the ability to form multiple different relationships with the same individual based on additional contextual information. In order to do this, though, the wolves would need to be able to perceive this combination of stimuli. The CSEM mechanism should be able to create these multiple relationships readily, once the perception system is able to glean the relevant conjunctions. Nevertheless, if there were multiple CSEMs for overlapping sets of contextual information, there ought to be some “bleeding” between them. The limitation here lies in the simplistic notion of context that is currently maintained by our system. A more elaborate model of context would allow for correspondingly complex relationships.

#### **7.1.3 Behavioral Simplicity**

A third issue arises from the fact that the virtual wolves’ behaviors are quite simple compared to those of real wolves. Increasing behavioral complexity could cause unforeseen ramifications for the CSEM mechanism and the AlphaWolf system as a whole. We address a few of these concerns in the Future Work section at the beginning of the next chapter.

#### **7.1.4 Relationship Bandwidth**

A clear limitation caused by all of the above issues and several others is the limited “social bandwidth” that may be transmitted by the AlphaWolf system, in particular with regard to relationships between two people at opposite ends.



**Figure 7-1: Limitations on social bandwidth are imposed by simplification and interfaces. The interaction bandwidth between the two people, for example, is greatly compressed by passing it through the virtual wolves (not drawn to scale).**

Each stage of the installation narrows the amount of social information being transmitted. When a person uses the Howl Interface to direct his or her pup, the full spectrum of possible human social experience is converted down to a few numbers. The social implications of these numbers are then further muddled as they pass through the virtual wolf and its behavior system. The final reduction comes because of the limited interaction bandwidth between wolves; there are only a few behaviors that each wolf can take towards the other. As Figure 7-1 shows, the narrowing of the “pipe” through which social relationship information must flow narrows drastically and then widens back out. The relationship between Human A and Human B that is facilitated by the installation can have no more dynamic range than the narrowest point, where the two wolves interact.

This drastic simplification is not necessarily a bad thing for a newly formed relationship among people, though. The relationships are simple and clear, and involve a very constrained set of possibilities. While it would certainly be intolerable to have this be the sole mechanism by which to interact with another person, it could serve as an easy, safe introduction, or as a novelty for an already-established relationship. In addition, people are skilled interpreters of social psychology; although the communication bandwidth of the actual relationship might be limited, the relevant psychology could be evoked by more simple information.

### 7.1.5 No Memories of Specific Events

Another limitation of the system is that it does allow a computational entity to capture memories of specific interactions. For example, if the gray pup habitually dominates the black pup, and then submits one time, the black pup will merely dilute its submissive relationship towards gray, rather than remembering that specific interaction. Although it is unclear whether or not animals form these kinds of specific memories, it is clear that people do.

One benefit of remembering a specific interaction is that the entity could try to learn what was special or different about that specific case. Returning to the example, perhaps the gray pup had just been trounced by the white pup. If the black pup were able to learn such things, it could learn to hang around and wait for white to trounce gray again, and “coat-tail” on that interaction. In the current social world of the virtual wolves, though, this level of complexity is well beyond those entities’ abilities, again, due to the system’s impoverished notion of context. Whether or not this

is how people learn is unclear as well. Perhaps we take mental snapshots of uncharacteristic interactions for some other purpose.

### **7.1.6 Real-time**

One of the central design decisions of the AlphaWolf system is that it run in real-time. This means that the wolves should always run at least 15-20 frames per second. However, because of the complexity of the wolves, their interactions, and the virtual world, the frame rate never gets much above 20Hz. This speed limit makes it difficult to perform long runs or multiple iterations of simulations with the system. As an example, the resource exploitation simulation data are derived from just a few runs, rather than from the averages of several thousand runs.

The real-time aspect of the project is one of its features, as well. It enables human interactivity, and allows the wolves to utilize the full Synthetic Characters architecture. Nevertheless, it is a limitation with regard to synthetic biology and artificial life applications, at least until increases in computational power render it irrelevant.

### **7.1.7 Supporting Technology Problems**

Subjects in the human user study wrote in a range of problems that they encountered while interacting with the installation. Most prominent among these problems were the poor collision avoidance, the sound effects, and the camera system. Each of these problems hurt subjects' ability to suspend their disbelief and become immersed in the experience. These three problems point out several limitations of the system as a whole.

First, it demonstrates that the installation is, as the saying goes, only as strong as its weakest link. If any aspect of the system is broken, the installation as a whole is greatly compromised. This phenomenon encompasses the entire environment in which the piece is presented; if the room were too hot, for example, that would have been a problem, too. The inter-reliance of all elements, both in terms of technology and in creating a single, unified experience, points to the importance of the whole-system approach that lies at the heart of the Synthetic Characters Group.

Second, it demonstrates the challenge of working in a domain that is full of unsolved issues. Real-time close contact among autonomous characters, interactive camera systems, and dynamic sound effect systems are all areas of active research. Making an installation that relies on unsolved technologies runs the risk of keeping people from becoming fully immersed.

Finally, it demonstrates that it can be a challenge to focus the attentions of an audience. In general, we try to leave out any installation elements that don't contribute to the central goal. The first time we learned this lesson was during the making of the *Swamped!* installation in 1997-98. Our animator made a gorgeous tree-house for the background of the virtual world – beautifully designed and texture-mapped. We all loved it, and immediately stuck it in the world. As soon as the first participants came to try out the installation, they immediately ran over to the tree-house, ignoring the rest of the installation, and becoming annoyed when they couldn't figure out how to climb into the tree-house. "We didn't implement it," we said. "Why not?" they asked. "It's not the main part of the installation." "Oh. Hrmph." Before long, we removed the tree-house, people began attending to the main area of the installation, and no one ever asked "Why don't you have a beautiful tree-house in the background?" If it's not there, they won't miss it.

### **7.1.8 Novelty Can Be a Problem**

Novelty has a major impact on users' experience of AlphaWolf. The user tests quantified this phenomenon, with people preferring the second of two similar runs by a wide margin. Later runs

proved to be more enjoyable, to have clearer relationships, to have more comfortable interfaces, to offer more control, and to provide a more immersing experience. The wolves in later runs acted more like real wolves. The lesson to be learned from this clear data is that if a designer wants someone to get beyond the interface and experience the “meat” of the installation, that designer needs to make sure that the initial experience is interesting enough and easy enough to master that participants don’t get bored or frustrated and give up. The initial “hook” needs to be compelling. Interestingly, the “hook” doesn’t necessarily need to be compelling for the same reasons as the ultimate goals of the project. For example, the beautiful graphics of AlphaWolf helped draw people in long enough for them to “see” the social behavior.

Another lesson, which will be addressed again in the section on educational applications later in this document, is that since people take a little while to get into it, it’s probably better to have some sort of introduction, rather than leaping right into the heart of whatever material is meant to be communicated. Don’t put all your good material right up front. People will still be coming to grips with the interface.

## **7.2 *Future Work***

As was pointed out in the Limitations section, there are a number of aspects of real wolf behavior that are not addressed by the current model. Many of these are clearly implementable already, while other would require a bit more thought.

### **7.2.1 *Scaling***

The AlphaWolf social relationship mechanism is the simplest version that I could create. One of the great strengths of a simple model is that it can be scaled and extended to create a more complex model. While the AlphaWolves have never existed in a pack of more than 6 wolves in their full, graphical form, it seems reasonable that the same mechanism could work in larger packs, or in a group of packs that occasionally interact and trade members. The existing model allows for the dynamic introduction of one or several new members into a pre-existing social order (for example, in the AlphaWolf installation each of the three young pups meets the adults in the virtual pack); therefore it should be able to allow many groups to negotiate their relationships even when members occasionally disperse to another pack. Nevertheless, it would be interesting to explore the range of parameters – group size, rate of dispersal, social learning rate, etc. – that create stable social orders rather than deteriorating into random interactions.

Sometimes large-scale systems have emergent properties that are not noticeable in smaller groups of entities. The analysis of multiple packs existing, interacting, foraging and dispersing, might yield interesting and unexpected results. It is reasonable to expect that the AlphaWolf mechanism could serve as the underpinnings for such a system because it is simple enough, computationally, that large numbers of virtual wolves might be able to be run. The two main ways these large-scale studies could be possible would be to run fully graphical wolves in a non-real-time fashion (since pack size is already at the limits of modern hardware’s performance), or to create wolves who do not have the computational overhead of full graphics and behavior systems. However, this latter case would lose quite a few of the attributes that make AlphaWolf so compelling – interactivity, visible expressiveness, etc. As hardware improves, though, it should become possible to run packs of wolves in real time on a much larger scale.

One problem that would need to be dealt with in a larger group of wolves is the topic of relationship culling. Currently, each wolf forms relationships with all the other wolves; since relatively few wolves exist in the world at the same time, it has been unnecessary for the wolves to “forget” about wolves whom they hadn’t seen for a while. However, if a wolf were to

encounter many other wolves during the course of its life, it would need some mechanism by which to cull those relationships that were no longer relevant.

### **7.2.2 Single-Sex Hierarchies**

One of the phenomena that appears to occur in captive wolf packs is the existence of separate hierarchies for the two sexes of wolves [Mech 1999]. As Mech says, the main use of dominance hierarchies in the wild is for food distribution [Mech 1999]. Dominance is therefore closely tied to a scarce resource.

When pups in the wild come of age to breed, they disperse and form their own packs. Therefore, there is little competition for mating opportunities in a wild pack. In captivity, however, where several adult wolves of each sex may be kept in the same relatively small enclosure, mating opportunities become a scarce resource, in particular as the breeding season approaches when all members of each sex are competing for one spot. The competitive groups are therefore different from the ones competing for food (where it's every wolf for itself, regardless of sex), and might command a second set of hierarchies.

There is no clear reason why the current AlphaWolf mechanism could not form the basis for an exploration of this topic. By causing each wolf to form one CSEM for each wolf around the topic of feeding, and another CSEM for each same-sex wolf around the topic of breeding, and allowing some interplay between the two kinds of CSEMs (for example, by having multiple CSEMs for the same individual bleed into each other), it might be possible to create two distinct but inter-relating hierarchies.

### **7.2.3 Deposed Alpha Plummetts**

A phenomenon that occurs in wild wolf packs, but that does not occur in the current model, is that deposed alpha-ranking individuals often plummet in dominance after they are deposed, sometimes becoming the lowest-ranking individual, or “scapegoat” of the pack [Klinghammer 2001]. There are several possible explanations for why this drastic rearrangement of the dominance structure happens. First, it is possible that the damped positive-feedback system that is the dominance hierarchy keeps a wolf in the top rank for a certain amount of time after its physical prowess has begun to deteriorate. Therefore, when the individual is finally deposed, it has further to fall than might be expected. This possibility could be represented in our system by tying dominance success into some changing physical ability. Currently, all wolves are created equal, as far as fighting ability goes, and are only differentiated by the CSEM mechanism itself.

A second possibility involves the prospect that wolves can perceive the dominance and confidence of another individual, and modify their behavior accordingly. While it is unclear if the ability to “smell fear” is strictly an olfactory skill, it is certainly the case that animals can detect the suites of characteristics that correspond to certain emotional states. How, then, could this be implemented in the AlphaWolf system? Currently, wolves base their behavioral decision to dominate or submit on their emotional state. If a wolf has a strong emotional memory of another wolf, that has a very strong impact on that emotional state. If, then, a wolf feels very submissive to another wolf in the current system, it is unlikely that it will decide to dominate that wolf, regardless of how that wolf behaves. Consider, instead, the possibility of each wolf having a global confidence, which is some function of its several most recent interactions. If this global confidence were perceivable by another wolf, that wolf could then use the net confidence of both itself and its current interaction partner to determine how much to apply its CSEM. In this model, it might be possible for a wolf who is “shaken up” (e.g., recently deposed) to lose its next

interaction *because* of its reduced global confidence, and the fact that the other wolves could see that it had “lost its nerve.”

A final possibility, with regard to implementing the plummeting of a deposed alpha, is that an alpha wolf is held in place to a certain extent by its alliances with the other members of the pack (although, as I mention below, there is no clear evidence that wolves form alliances). Having an alliance with the most dominant individual would be of significant value to a subordinate wolf, with regard to resource access. Once an individual is no longer in the alpha position, though, it might lose its alliances (when most of the pack shifts their allegiance to the new alpha), and would therefore lose much of its social “clout.” This dramatic shift, especially if coupled with the deterioration of actual strength mentioned above, could account for the “scapegoating” of a deposed alpha. Because of the lack of clear alliances in wolves, this technique might be more appropriate for a simulation of dominance among primates.

#### **7.2.4 Coopting Behaviors to Serve as Social Signals**

In wolves, submissive behaviors appear to be derived from infantile behaviors. The adult submissive display described above by Allen bears a strong similarity to the reflex urination found in very young pups.<sup>12</sup> Fox [Fox 1971] refers to social behaviors derived from infant behaviors as socio-infantile patterns. “The submissive activity is, in its essence, an activity of the cub.” [Schenkel 1967] Currently, the virtual wolves perform submissive behaviors because they have been hard-coded into their behavioral repertoire as the appropriate behaviors to perform when one’s dominance value is low.

The AlphaWolf system could serve as part of the basis for an implementation of the above phenomenon. If adults, when disciplining a pup, would not relent until the pup offered a characteristic submissive behavior, which is already in their repertoire for a different purpose, those pups could be made to learn to perform that behavior in the novel context of being disciplined. Only by having an expectation of an ensuing interaction, though, (e.g., “that approaching adult is probably going to discipline me”) could a pup learn to submit appropriately before the adult actually arrived. Real wolf pups do that. The AlphaWolf system could be the mechanism by which those expectations were formed.

#### **7.2.5 Facilitating Social Learning**

In fact, the above example of learning to give social signals in the presence of a disciplining adult is just one example of a larger suite of phenomena known collectively as social learning. Many instances of social learning rely on an individual having some expectation of the interaction that will occur with another individual; the AlphaWolf system could help enable this in simulation.

For a full treatment of social learning, the reader is directed to the book “Social Learning in Animals” [Heyes 1996].

#### **7.2.6 Signaling, Deceiving and Detecting Deception**

In natural systems, the process of communication creates the possibility for deception. As I mentioned in the Related Work chapter above, the ability to remember a specific individual could serve as a mechanism for biasing one’s behavior against deceivers in the future. By increasing the cost associated with deception, deceptive behavior should be reduced.

---

<sup>12</sup> See the Related Work section for more information.



In the AlphaWolf system, there is no mechanism by which deception could take place. The wolves maintain honest models of their interactions, and update them consistently based on their interactions. Nevertheless, it would be interesting to use the AlphaWolf system as a platform for exploring deceptive communication. In order for such experiments to be possible, there would need to be actual costs associated with fighting, and a signaling stage that was distinct from fighting. By allowing an individual to assess its own ability and the relative ability of another, as well as the value of the resource in question, the system could create computational entities who could obey a variety of strategies – honest signaling or varying degrees of cheating. The character-based system, which treats each character as a discrete unit, would be appropriate for simulating a competitive system in which each character knows more about itself than about its competitors. If deception could take place in the system and entities could detect that deception, those entities could form emotional memories of the level of deception remembered for each social partner.

### 7.2.7 Alliance Formation

While there is no incontrovertible evidence that wolves form alliances, there are several reasons to propose that they may. First, the seminar leader at Wolf Hollow Ipswich believed that two brothers had been “working as a team” to jointly climb their way up the pack’s hierarchy [Morss 2000]. (Nevertheless, Hemelrijk [Hemelrijk 1999] has suggested that apparent alliances can seem to exist without any internal representation for them.) Second, parents in the wild appear to endow their youngest offspring with the right to feed before their older siblings [Mech 1999]. While this phenomenon is not an alliance in the traditional sense of the term (since the pup can’t reciprocate, except by passing on its parents’ genes), it does appear to be a case of one wolf behaving differently towards a second because of the presence of a third (the parent keeps the older sibling from feeding because the pup is still eating).

In addition, other social animals (e.g., chimpanzees, humans) do form alliances [Harcourt 1992]. Therefore, it is a very interesting to consider how the AlphaWolf mechanism could be used as the basis for the formation of alliances.

The way in which alliances could be formed is to take advantage of one of the basic premises of the CSEM mechanism. As I mentioned in the introduction, a CSEM is a *Context* Specific Emotional Memory, rather than a *Partner* Specific Emotional Memory because it has always been our intention to extend this work to allow emotional memories to form at an arbitrarily fine or coarse grain. Forming a CSEM about Wolf A in the presence of Wolf B, and a separate CSEM about Wolf A *not* in the presence of Wolf B would be a simple way to represent an alliance within the AlphaWolf structure.

One potential drawback of this increase in complexity is that it scales much less well than the simpler model. While the complexity of the simple model scales geometrically, the alliance-forming model scaled exponentially. However, this scaling problem may account for the clear limitations in group size found among social animals, which appear to correlate loosely with brain size [Dunbar 1993].

### 7.2.8 Conjunctions

The alliance formation example above points to a larger-scale problem – that of having multiple CSEMs for the same individual, depending on some additional contextual information. The notion of giving each wolf a CSEM for each other wolf in the presence of each third wolf is one case; others could include other attributes, such as Wolf A “when it smells like meat” or “when it’s near the den.” These potential conjunctions represent a significant challenge for the extension

of the AlphaWolf mechanism into the domain of relationships with human-like complexity. Having a computational entity that can figure out which are the salient bits of contextual information to use when forming emotional memories, and can dynamically augment and cull its collection of CSEMs, is a major and important area of future work.

### **7.2.9 More Complex Emotional Models**

Another area of extension for the AlphaWolf system is in its emotional model. Currently, the CSEMs formed by the wolves apply to dominance only. However, the system was built with the intention that it could eventually be extended to use the full Pleasure-Arousal-Dominance model proposed by Mehrabian and Russell [Mehrabian 1974]. Having different CSEMs for each of the three axes would allow much more complex emotional interactions to occur.

In addition, the CSEM mechanism could be used with different emotional models. Both categorical and dimensional models (see section 3.1.3) could be used with the AlphaWolf system. Allowing events in the world, and in particular social interactions, to influence the emotions of the characters (e.g., surprising behavior leading to increased arousal), would insure that the emotional model was fully integrated into the scenario.

In order for a more complex models to be exhibited through the behavior of characters such as the virtual wolves, those characters would need a correspondingly complex set of expressive animations. In making AlphaWolf, the creation of the source animations was one of the rate-limiting steps; to make a wolf that could show a clear dynamic range in pleasure and arousal, in addition to dominance, would have entailed significantly more work for the animators.

### **7.2.10 Formalization of Development**

The way in which AlphaWolf addressed behavioral development was fairly *ad hoc*, with each ActionTuple having one or more TriggerContexts and DoUntilContexts pertaining to some hard-coded age-range. For development to play a more significant role in our system, we need a formal way of specifying groups of behaviors that are innate and groups that belong to a given critical period. It is important that the development mechanism integrate cleanly with the learning component of the system so that behaviors may be maintained through learning even after their critical period has elapsed. This combination of development and learning lies at the heart of the socio-infantile patterns described by Fox [Fox 1971].

A possible implementation of behavioral development would entail adapting the Synthetic Characters Group's action selection mechanism. Currently, 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. I propose adding a fifth optional component called an AgeContext. An AgeContext will have an "Age of Onset", an "Age of Offset" and a "Value." When the age of the entity is between the Age of Onset and the Age of Offset, the AgeContext's Value will be added to the TriggerContext. By integrating development with the action mechanism that is already in place, the entity's learning of developmentally timed behaviors will happen through our existing learning mechanism.

### **7.2.11 Synthetic Sensing for Assessing Dominance**

One topic that our group has felt would be very exciting to implement is the ability for the wolves to use synthetic vision for detecting dominance. One of the criteria for our expressive ranges is that they be perceivable by people because of the quality of the animation; perhaps it would be possible to make a vision system for the wolves that would view their social partners' behaviors

and thereby determine their motivational state. Currently, the perception model that the wolves use is far simpler than this synthetic vision would be.

For example, this system might be able to learn that tail up and ears up correlate with dominance, or that certain other consistent patterns are only shown by submissive wolves. The fact that subject were so successful at decoding the social relationships among the pups (see section 5.2.3.1) demonstrates that there were certain clear patterns in the wolves' behavior. Having a system that could detect the same traits that people can evidently detect would open up a diverse range of potential signaling and communication processes that real wolves use that we have yet to explore in our system.

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 aggressively 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.

### 7.2.12 Close Contact

The biggest technical problem that participants in the human user studies noted was the poor collision detection. Avoiding collisions among a group of several mobile and unpredictable entities is quite difficult, especially with a limited repertoire of behaviors and no computational model of physics. However, in order for virtual wolves to exhibit the kind of vigorous social behavior that real wolves habitually exhibit (see Figure 7-2), it will be necessary to address a variety of hard problems in motor control, collision detection and avoidance, and synthetic physics.



**Figure 7-2: Close contact – one of the reasons that social behavior is a hard problem.**

The avoidance of obvious errors such as interpenetrating wolves is crucial for participants' suspension of disbelief. While people can evidently tolerate quite a wide range of behaviors from the virtual characters without losing their feeling of immersion and realism, as soon as two pups occupy the same physical space, people are yanked back into the real world, where they are

watching a (broken) simulation on a computer screen in a lab somewhere. Nevertheless, it is a testament to the success of AlphaWolf, or to the tolerance of the participants, that the interpenetrations did not prove to be a “show stopper” with regard to their enjoyment of the installation.

## **7.3 Applications**

The above sections have mentioned a variety of ways to extend the AlphaWolf pack in more-or-less its current form. The broad ideas in this thesis, though – about computational entities who can form social relationships with each other and with people – have applications in a wide variety of commercial and academic pursuits. This section attempts to give a cross-sample of possible ways in which social computational entities could find their way into the world around us, some of which could come about in the next few years, and others that are a bit longer-term.

### **7.3.1 Entertainment**

The entertainment industry has perhaps the most pressing need for social synthetic characters. Computer games, television, movies, toys and interactive theme park installations could all benefit from a mechanism to make social computational entities.

#### **7.3.1.1 Computer Games**

The computer game industry needs synthetic social relationships for their good guys, bad guys and semi-autonomous avatars. The virtual worlds of future video games will be populated with convincing characters featuring elaborate mechanisms for simulating social competence. Characters who can remember the players and each other, and form friendships or adversarial relationships, will provide a much more exciting backdrop for a wide assortment of game genres.

In addition to helping create new characters who can engage players and each other in social ways, the research described in this thesis could help make interactive versions of existing non-interactive characters. 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 these entities be controlled by a user and yet stay “in character”?

The challenge of building directable virtual wolves who nevertheless exhibit plausible wolf behavior is very 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?” The clear division between action and emotion in AlphaWolf has proven to be a useful mechanism for making semi-autonomous characters who obey the direction of a human participant and still present a consistent personality. This division could help make consistent and yet directable themed characters.

#### **7.3.1.2 Television and Movies**

Social computational entities could also serve a creative role in the television and movie industries. In movies, computer graphical crowds, flocks and herds (e.g., *Antz*, *Jurassic Park*, *Lord of the Rings*) already exist. As mechanisms for synthetic social behavior get better, synthespians will take on leading roles without relying as heavily on hand-animation. Directing virtual actors in the same way that one might direct human actors – that is, using social roles, emotions, and motivations – will some day (soon) be more efficient than hand-animating those

virtual actors. To make this kind of direction possible, the virtual actors will need to know how social and emotional phenomena connect with each other.

In the area of animatronics, too, social competence could be of use. As Stan Winston, founder of Stan Winston Studio, pointed out in a recent lecture at MIT, real-time expressiveness in animatronic characters (currently done by puppeteers) is necessary for scenes to “come alive” for the human actors who interact with them. Actors know that their best performances come when they are doing a scene with another skilled actor; therefore, doing an emotional scene to a blue-screen is one of the hardest dramatic challenges, and often falls flat. That is why Stan Winston’s creatures, who actually exist on stage with the actors, not only look better than a computer graphical version of the same thing, but also elicit better performances from the human actors in the scene. Synthetic social behavior could help automate these mechanical actors, giving them the scaffolding they need to begin to understand why they should act one way or another.

### **7.3.1.3 Toys**

Toys that can form relationships with each other and with the children who play with them also offer interesting possibilities. For example, a toy that could tell the difference between a child and his friends could play with them differently based on the history of play patterns between them. Or consider a set of dominoes that, rather than just falling down and knocking over whichever other domino is in the way, could run to a specific other domino and knock it over. Or imagine a crowd of plastic cave men who learn to be scared of a plastic saber-tooth tiger over the course of a play session. Although scenarios like these might sound like the opening scene to a Hollywood horror movie, I imagine that social toys will be much less malevolent, simply making playtime more stimulating and fun for children.

### **7.3.1.4 Location-Based Entertainment**

Interactive installations, arcades, amusement parks and other location-based entertainment could also be improved by characters with social and emotional abilities. For many of the same reasons that computer games and movies will embrace social synthetic characters, these forms will as well. Disneyland, Magic Mountain and other major sites are attached to specific characters; making interactive versions of these characters to populate the various rides would make those experiences better and reaffirm the branding of Mickey and Bugs.

### **7.3.1.5 New Forms**

Perhaps the most powerful argument for the value of social synthetic characters in entertainment comes from the experiences that we can’t conceive of yet. To date, there *haven’t* been computational entities who could form social relationships; therefore, trying to shoe-horn them into existing media, while possible, probably isn’t the best use of them.

Over the last few years of late-night conversations, we’ve thought about a lot of possibilities for our characters. One of the metrics I use when evaluating a potential experience is the “first date test.” Would the experience make a good first shared event for a person and his or her potential significant other? Some of the criteria for a good date are:

- Safe – when two people don’t know each other well, it’s better to do something relatively public.
- Fun – having a good time together creates a common bond.
- Interesting – gives you something to talk about over dinner or coffee.
- Conspiratorial – both people should preferably be “on the same side” of the experience.

Movies clearly satisfy all of these criteria. Going to a theater with a few hundred other people is reliable protection against something going terribly awry. Movies are (usually) fun. The plot and characters of a movie give the budding couple a topic of conversation to fill any uncomfortable silences. And, since most movies have a built-in bias for or against certain characters (even bank robbers can be protagonists), the couple is brought together by a joint favoring of certain characters (which can be pleasantly juxtaposed with personal preferences, as appropriate).

For an example of the kind of experience that doesn't exist yet but will nevertheless benefit from computational characters with social relationships, consider the domain of interactive film. Movie theaters are gradually shifting from celluloid to digital projectors; as more theaters "go digital", there will be fewer reasons for the media that they play to be completely fixed. Rather than playing a DVD into the projector, there could just as easily be a computer playing out a subtly interactive version of the same thing.

As an example of how this interactivity could enhance the subjective experience and commercial viability of a movie, consider the following possible future movie-like entertainment experience. Upon coming to the theater to watch "Batman X", each viewer calls a certain phone number with their cell phone, to "sign in" at the theater. The movie comes on, and everyone watches the same events unfold on screen. At some point in the film, though, each viewer's cell phone vibrates, and he or she is asked a question. "Press one if it's raining tomorrow in Gotham City. Press two if it's sunny." The viewer whispers with his friend, perhaps, and they decide to make it rain. Each viewer gets to control one tiny bit of interactivity in that particular run of the movie. The interactive movie adapts to incorporate these minor variations in setting. The plot doesn't change significantly; Batman still catches the crook in the end. But there are several main areas in which value has been added to the experience:

- Each viewer feels a bit of "ownership" over that particular viewing.
- Each theater gets to see a slightly different version of the movie, but the plot was close enough that people from different theaters can still talk about it over the water cooler at work the next day, and also talk about the differences.
- Different theaters offer a different experience, due to the composition of the audience. While one theater takes the more sinister of the two options ("Yes, it's raining."), the other gives a consistently happier version.
- Every time someone watches the movie, whether at the same theater or different, the movie changes subtly, so that multiple viewings offer a more thorough understanding of the lead characters and their motivations.
- People want to come back again and again, and see a movie in different settings, thereby delighting the movie's producers and the theater owners with the increased business.

The trade-off for these benefits is that, instead of being a few thousand feet of celluloid, the movie must be a robust, adaptable, interactive program running convincing synthetic actors and beautiful real-time graphics. AlphaWolf is a first step down the road to that goal. Plus, an interactive movie would make an interesting first date.

### **7.3.2 Education**

Education is another significant area in which social synthetic characters might be applied. Social toys could be used to help teach children that their actions have long-term effects on the social entities around them – people, pets, etc. These toys could adapt their level of social complexity to the skill level of each child. In addition to helping to teach social skills, socially enabled virtual instructors are already being developed for a variety of topics. Finally, social

characters could be used to teach kids about social dynamics in other species. Imagine a virtual laboratory at a zoo, populated by computational animals, where visitors could set various parameters on the animals and watch as their social structures changed from bird-like to wolf-like to chimpanzee-like.

### 7.3.3 Robotics

Throughout this document, I have returned occasionally to various robotic applications for computational entities with social relationships. Robotics has been the main real-world application of the machine-machine kind of relationship that is enabled by the AlphaWolf mechanism. As was described in Chapter 5, multi-robot systems could use social relationship mechanisms like the one in the AlphaWolves to negotiate their interactions. In addition to facilitating the interactions of multi-robot systems, interaction paradigms inspired by social animals could make those systems easier for people to understand and control. For a fuller treatment of these topics, the reader is directed to [Arkin 1998].

Vaughan et al. hypothesized that “a dominance hierarchy will be effective only when there are (i) non-uniform abilities in the group and (ii) a relatively slow change in the abilities of individuals. These conditions are relatively unusual in robotics with the important exception of systems with learning, evolution or other long-term adaptation.” [Vaughan 2000, p. 9] That is exactly the reason that we find social relationships, and in particular emergent dominance hierarchies, so exciting, because the work of the Synthetic Characters Group focuses on learning.

There have been vast ranks of science fiction books, movies and television shows that feature social robots in one form or another – R2D2 and C3PO from *Star Wars*, Commander Data from *Star Trek: the Next Generation*, many of the robots in Asimov’s stories (e.g., Robbie in [Asimov 1950]), *The Terminator*, the androids in the *Alien* series, David and Teddy from *A.I.: Artificial Intelligence*, etc. While not all of these are the most benevolent of socially enabled robots, they do point to the wide range of potential applications for a computational social relationship mechanism. (And until these future possibilities get made in reality, social computational systems can also help bring them to the silver screen.)

To provide two specific examples, I will describe how the AlphaWolf mechanism might function in a robotic social companion and in a robot administrative assistant.

Imagine a retirement community with a robotic pet. The pet should be largely autonomous, wandering around to interact with different members of the community. It should be able to form memories of interaction styles that individuals prefer; one woman likes to have it sit on her lap and purr, another gentleman prefers a rousing game of fetch. In addition to having the autonomy to remember preferred behavioral patterns and styles, the pet should be directable by the staff of the facility. Perhaps one woman recently lost her husband and needs particular attention for a number of weeks. The robot would have no way of knowing that much about the social context, but could nevertheless be directed to attend to her preferentially for a time. The interface might be as simple as bringing it on a leash to visit the woman; thereafter, it has a preference for her until it is redirected or its memory fades. To set the rate of fade, however, might entail a much different kind of control – a more traditional computational interface, for example. In addition, it could be appropriate to place certain limits on the pet’s exuberance around more frail individuals.

As another example, consider a robotic administrative assistant. The assistant might try to learn about different circumstances, model its boss’s emotional state, and model the relationships between its boss and the boss’s co-workers. If the boss’s emotional state seemed tense, or if they

were interacting with a co-worker with whom the assistant was unfamiliar, it might behave conservatively, guaranteeing that it performed basic tasks as expected. If, on the other hand, the boss seemed relaxed and with friends, it might be a little more experimental, suggesting projects that had been back-burnered at some point in the past. In addition to the adaptive elements, however, the assistant should maintain a core of functionality that responds promptly and accurately under all circumstances, and should also have a direct interface by which the boss could specify or correct the assistant's behavior.

Both of these examples point to a dual interface for a real-world social entity – one, a social interface that treats the robot as an independent entity, and the other, a “super user” mode that allows direct authorship of the character's character. In fact, that is very much how the AlphaWolf system is designed. On the one hand, there is the Howl Interface through which people can interact with the wolves. On the other hand, there is a set of authorial tools with which we design and modify the wolves. The nature of the interface depends greatly on the context in which the interface will be utilized.

### **7.3.4 E-business**

Social structures that traditionally have been handled by people are becoming increasingly codified as software agents. Automatic collaborative filtering systems [Shardanand 1995] and expert finder systems [Lieberman 2000], for example, are serving in roles that used to be primarily the domain of friends and associates. E-commerce agents are doing financial work that used to be done by humans. As agents take over more complex relationships from people, it will become necessary to integrate the hierarchical social knowledge that people use in everyday life – ideas of social dominance and submission. We all use and understand the subtle cues that inform us of where we stand with respect to each other (e.g. body language while haggling). We all make alliances to achieve mutual benefit (e.g. buyers' collectives). Soon, our agents will too.

There are several main ways in which the themes of this research can be applied to multi-agent systems. Dominance hierarchies can be used to streamline negotiations among agents who have already established a relationship in the past. Social status can assist in the formation of alliances among agents who might mutually benefit from some kind of collaboration. Finally, an appropriate representation of dominance and submission can help humans interact with a system.

Businesses have many analogs of these social hierarchies, especially with regard to alliance formation. Just as two adult male chimpanzees may team up on a third individual, more powerful than either, and defeat him, two or more companies often bundle their products together to provide added value to the consumer and therefore win greater market share. For example, Microsoft and Intel dominated their market during much of the 1990s as a result of their alliance. “The sheer muscle of the so-called Wintel (Windows-Intel) combine, analysts say, has kept any other computer architecture or operating platform from thriving to a degree that would even remotely threaten its dominance.” [TCSGlobalNews 1996]

In the business world, as in the natural systems described above, there is an expensive switching cost associated with overturning the dominance hierarchy. Whereas a wolf has to spend energy and runs the risk of injury, a company that decides to switch operating systems has to put in quite a bit of effort and runs the risk of significant problems in the process. This acts to keep a given dominance hierarchy in place, rather than encouraging dithering between two closely matched alternatives. The high cost of switching is enough to outweigh the small advantage to be gained by a product that is only marginally more valuable to the consumer.



In both animals and businesses, the sheer fact of being dominant makes an individual a valuable alliance partner. Therefore, dominance tends to be self-reinforcing. This, too, leads to a preservation of the current structure.

### **7.3.5 Synthetic Biology and Artificial Life**

Synthetic social behavior is relevant to a variety of scientific research and engineering tasks. The field of animal behavior, which already utilizes a range of computational modeling techniques, could continue to benefit from work in the modeling of social behavior.

An interesting experiment that could arise from 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.


In addition, the study of artificial life using full 3D-graphical virtual creatures could offer some results that might not be achievable without the graphical representation. As an example, the ability to tell when one wolf pup was looking at another wolf pup was crucial to debugging the interaction model among the pups. As artificial life begins to model creatures that are more and more complex, having a graphical front-end and expressive social abilities might be seen as more than just “window dressing.”

### **7.3.6 Fine Art**

While perhaps not as commercially viable as some of the other applications described here, fine art could benefit from social computational entities. The work of Marc Downie, a graduate student in the Synthetic Characters Group, points to this application. Downie’s projects, for example “music creatures” [Downie 2001b] and “Loops” [Cunningham 2001], often include multiple interacting virtual entities. The way in which these entities interact reflect ideas of communication and social interaction; his characters change their behavior toward each other based on their interaction history.

It is my hope that AlphaWolf can be seen as an artistic endeavor as well as a technical showcase. If that is the case, perhaps its use of computational entities that form social relationships could inspire other art projects involving social groups of computer systems. The recent awarding to AlphaWolf of an Honorary Mention in the Prix Ars Electronica provides some support for the premise that synthetic social relationships can form the basis for technological art.

In addition, interactive art works could be made possible by the directable aspect of the social relationship mechanism described here (see Figure 7-3). While computer-enabled interactivity is still in its infancy as an artistic medium, it could hold promise for future artistic efforts.



**Figure 7-3: The author, about to awaken a sleeping virtual wolf pup.**

### **7.3.7 Computer-Mediated Communication Technologies**

AlphaWolf helped create relationships among the people who interacted with it; these relationships point to a possible application of the social relationship mechanism in the area of computer-mediated communication technologies. Various technologies take an active role in the mediation of interpersonal communication. As Judith Donath points out, “in order to foster the development of vibrant and viable online communities, the environment - i.e. the technical infrastructure and user interface - must provide the means to communicate social cues and information: the participants must be able to perceive the social patterns of activity and affiliation and the community must be able to evolve a fluid and subtle cultural vocabulary.” [Donath 1996] A social relationship mechanism like the one in the AlphaWolf system, and socially competent virtual characters like the wolves, might be able to contribute to this vision.

Modern avatar-based chat rooms are a prime example of why these technologies are necessary. Social signals are one of the key aspects of a person’s behavior that he or she might want to communicate; nevertheless, social signals are challenging to specify in real time. An avatar or other representation of the self that has the ability to generate social signals that are appropriate to the relationship that the person has with his or her current interaction partner would offer a significant increase in communicative power and ease of use. Just as AlphaWolf allowed people to take on distinct roles with respect to each other, new kinds of mediating technologies could assist people in declaring their social intentions.

### **7.3.8 Human Computer Interaction**

When people interact with each other, we obey a complex set of social protocols in order to facilitate our interactions. When computers interact with each other, they obey a complex set of technical protocols to facilitate their interactions. If people and computers mean to interact in any meaningful way, either we need to learn to talk TCP/IP, or computers need to come to grips with what it means to be social.

On a long time scale, the broad topic of human-computer interaction could benefit in a variety of ways from the capability of computational entities to form relationships. Humans are inherently social entities; the main interaction paradigm that we use when interacting with complex systems

is the set of social rules. Our machines should be able to take advantage of this in-built predilection. This long-term area of application pertains to many of the topics discussed above, and to an assortment of potential future computational entities.

### **7.3.9 Artificial Intelligence**

As a final thought on the application of social relationships to computational systems, consider the Machiavellian Intelligence Hypothesis [Byrne 1988], which proposes that human-level intelligence arose as a way for us to maintain our elaborate web of social relationships. Perhaps if computational entities had to keep track of social partners, they might be on their way to “doing things that would be called intelligent if done by men.” [Minsky 1968]

## 8 Conclusions

This document has presented a system through which computational entities can form social relationships with each other and with people. The AlphaWolf system creates a virtual wolf pack in which the wolves form dynamic social relationships with each other; a major installation, a series of human user studies and a set of simulations built with this system confirm that people are willing to accept and participate in these relationships.

The Context Specific Emotional Memory (CSEM) mechanism that lies at the core of this system is a simple, robust representation of a social relationship. In the AlphaWolf installation, the user tests, and the resource exploitation simulations, the CSEM mechanism and its supporting technologies have been used to create an accurate, though simplified, sketch of wolf social behavior.

The essential function of the CSEM mechanism is to allow entities to learn associations between specific individuals and specific emotional states. When the entities meet again, their emotional memories kick in and cause them to return to a similar emotional state. The evaluations described in Chapter 5 prove that the AlphaWolf system works; it creates relationships that are recognizable by people, and that have an effect on people when they interact with them.

The necessary elements of the AlphaWolf system, as confirmed by the evaluations performed in that chapter, include emotion, perception and learning. Without any one of these components, the mechanism loses some element of its functionality (see section 6.1.1).

The relationships created by the AlphaWolf mechanism are explicitly designed with human interaction in mind. Human participation in and direction of the relationships formed by the AlphaWolf mechanism is central to its functionality. In the case of the AlphaWolf installation, the actions of the semi-autonomous pups were directed by human participants, while the pups themselves were maintaining their social relationships and expressing them through the emotional style in which they took the directed actions. In the resource exploitation simulations, on the other hand, the dominance hierarchy was directly influenced by causing a certain individual to be dominant in all of its interactions. The theme of directability has featured prominently throughout this document.

Synthetic social relationships have three main areas of usefulness, in accordance with three kinds of relationships that they help form. The machine-machine relationships have a variety of applications in multi-robot and multi-agent systems, some derived from the benefits of natural social systems, and others that are unique to groups of computational entities. The human-machine relationships that are formed could be valuable to a variety of human-computer interface domains, from entertainment to education to directable synthetic social groups. Finally, computational representations of social relationships could help mediate relationships among humans, for a variety of communicative goals. These applications were discussed in full in Chapter 7.

The wide range of possible applications of synthetic social relationships demonstrates that the research presented here is a significant contribution to a number of fields. Making computational systems take advantage of sociality both in their own interactions and in their interactions with people has the potential to improve greatly the efficiency and usefulness of those computational systems. A simple, robust representation of social relationships is an important step towards socially competent computational entities.

This work is derived from a variety of research in several different disciplines. For example, it draws on the biology of the gray wolf (e.g., [Fox 1971; Klinghammer 1985; Mech 1998]), Damasio's Somatic Marker Hypothesis [Damasio 1994], Mehrabian and Russell's emotional model [Mehrabian 1974], research in social computational systems (e.g. [Reynolds 1987; Breazeal 2000]), and previous work by the Synthetic Characters Group (e.g., [Blumberg 2001 (to appear)], [Burke 2001; Isla 2001b]). Further related work was described in Chapter 3.

While inspired by these previous works, the central ideas in this dissertation are original. AlphaWolf is the first system to learn and use emotional memories for social relationships in a real-time interactive environment. The real-time nature of the project is crucial to the system's value, in that it allows people to participate in and direct the relationships (see section 6.2.9).

This dissertation, and the body of research that it represents, confirms the central hypothesis that the AlphaWolf mechanism is a simple, robust, social relationship mechanism that can be used as the basis for a variety of relationships among machines and people.



**Figure 8-1: The black wolf howls to find its pack mates.**

Despite the relative success of this research project, there is still a long way to go before computational entities form social relationships on par with those of a real wolf or human. Chapter 7 describes three areas of future work that could follow on the work described here. First, there are issues of scalability – making groups of social computational entities that number more than six, and exhibit more complex behaviors than those of the wolves. Second, there are issues in extensibility – additions to the mechanism itself that would cause the social behavior of the virtual wolf pack to resemble more closely the social behavior of real wolves, or other real social systems. Finally, there are applications – long term projects that could draw on the ideas described in this dissertation. Without work in each of these three areas, computational entities with the kinds of social relationships described here will not have a significant impact on the world around us.

There are a number of philosophical issues that have not been addressed in this document. For example, there is a certain science-fiction element to the topic of social computational entities. The various ramifications of social and sentient machines are well-trod turf in written fiction (e.g., [Asimov 1950; Conklin 1954]), movies (e.g., *2001: A Space Odyssey*, *The Terminator*, *AI: Artificial Intelligence*) and television (e.g. *Star Trek: The Next Generation*). The author himself has even published on the topic [Tomlinson 2000a]. On the issue of whether or not sentient social machines should be caused to exist, I'll simply say this: People around the world create beings every day who are stronger, smarter, and more capable than they are, and yet it does not make headlines, and they do not feel threatened. Having children is part of every culture. If the creation of technological entities can be seen as a process closer to raising children than to building bombs, we can enjoy the rapid advances of technology without the fear that traditionally accompanies it. Making machines into social entities will help that happen.

## Bibliography

- Adams, E. S., and Mesterton-Gibbons, M. (1995). "The cost of threat displays and the stability of deceptive communication." Journal of Theoretical Biology **175**: 405-421.
- Alcock, J. (1989). Animal Behavior. Sunderland, MA, Sinauer Associates, Inc.
- Allison, D., Wills, B., Hodges, L., and Wineman, J. (1996). Gorillas in the Bits, Georgia Institute of Technology: GIT-GVU-96-16.
- Arkin, R. (1998). Behavior-Based Robotics. Cambridge, MA, MIT Press.
- Asimov, I. (1950). I, Robot. New York, Doubleday.
- Assanie, M. (2002). Directable Synthetic Characters. AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment.
- Axelrod, R. (1984). The Evolution of Cooperation. New York, Basic Books.
- Badler, N., and Allbeck, J. (2001). Towards behavioral consistency in animated agents. Deformable Avatars. N. Magnenat-Thalmann, and Thalmann, D. Dordrecht, Kluwer Academic Publishers.
- Balkenius, C., and Moren, J. (2000). "A Computational Model of Context Processing." Simulation of Adaptive Behavior '00: 256-265.
- Ballard, D. (1997). An Introduction to Natural Computation. Cambridge, MA, MIT Press.
- Barrett, L., Henzi, S. P., Weingrill, T., Lycett, J. E., and Hill, R. A. (2000). "Female baboons do not raise the stakes but they give as good as they get." Animal Behaviour **59**(763-770).
- 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.
- Beaver, B. V. (1999). Canine Behavior: A Guide for Veterinarians. Philadelphia, W. B. Saunders Company.
- Bekoff, M. (1974). "Social play in coyotes, wolves, and dogs." Bioscience **24**: 225-230.
- Blumberg, B. (1998). Swamped! Using plush toys to direct autonomous animated characters. Proceedings of SIGGRAPH 98: conference abstracts and applications.
- Blumberg, B. (1999). (void\*): A Cast of Characters. Proceedings of SIGGRAPH 99: conference abstracts and applications.
- Blumberg, B. (2001 (to appear)). D-Learning: What learning in dogs tells us about building characters that learn what they ought to learn. Distinguished Presentations in AI. G. Lakemeyer, and Nebel, Bernhard. San Francisco, Morgan Kaufmann.

- Blumberg, B., and Galyean, T. (1995). Multi-Level Direction of Autonomous Creatures for Real-Time Virtual Environments. Proceedings of the ACM Computer Graphics, SIGGRAPH 95.
- Blumberg, B., Tomlinson, B., and Downie, M. (2001). Multiple Conceptions of Character-Based Interactive Installations. Computer Graphics International.
- Blumberg, B. M. (1996). Old Tricks, New Dogs: Ethology and Interactive Creatures. Media Laboratory. Cambridge, MA, MIT.
- Bonabeau, E., Dorigo, M., and Theraulaz, G. (2000). Swarm Intelligence: From Natural to Artificial Systems. Santa Fe, NM, Santa Fe Institute.
- Bond, A. H. (2000). Modeling social relationship by mutual commitment, mutual control and mutual regulation. Proceedings of the AAAI Fall Symposium 2000: Socially intelligent agents - the human in the loop.
- Brand, M. (1999). Voice Puppetry. Proceedings of the ACM Computer Graphics, SIGGRAPH 99.
- Breazeal, C. (2000). Sociable Machines: Expressive Social Exchange Between Robot and Human. Artificial Intelligence Laboratory. Cambridge, MA, MIT.
- Brooks, R. A. (1990). "Elephants Don't Play Chess." Robotics and Autonomous Systems 6.
- Brothers, L. (1997). Friday's Footprint: How Society Shapes the Human Mind. Oxford, Oxford University Press.
- Buck, R. (1984). The Communication of Emotion. New York, Guilford Press.
- 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.
- Busch, R. (1998). The Wolf Almanac. New York, The Lyons Press.
- Byrne, R. W., and Whiten, A., eds. (1988). Machiavellian Intelligence. Oxford, Clarendon Press.
- Cañamero, D. (1997). Modeling Motivations and Emotions as a Basis for Intelligent Behavior. Proceedings of 1st Int'l. Conf. on Autonomous Agents (Agents '97).
- Cassell, J., Bickmore, T., Billingham, M., Campbell, L., Chang, K., Vilhjálms, 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 the ACM Computer Graphics, SIGGRAPH 94.



- Cassell, J., Vilhjálmsdóttir, H., and Bickmore, T. (2001). BEAT: The Behavior Expression Animation Toolkit. Proceedings of the ACM Computer Graphics, SIGGRAPH 01.
- Chen, L., Bechkoum, K., and Clapworthy, G. (2001). A logical approach to high-level agent control. Proceedings of the 5th International Conference on Autonomous Agents.
- Chi, D., Costa, M., Zhao, L., and Badler, N. (2000). The EMOTE Model for Effort and Shape. Proceedings of the ACM Computer Graphics, SIGGRAPH 00.
- Cohen, M. D., Riolo, R.L., and Axelrod, R. (1999). 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.
- Conklin, G., (ed.) (1954). Thinking Machines. New York, Bantam Books.
- Coppard, K. (1999). Wolves. Edison, NJ, Chartwell Books, Inc.
- Coppinger, R., and Coppinger, L. (2001). Dogs. New York, Scribner.
- Coppinger, R., and Schneider, R. (1995). Evolution of working dogs. The Domestic Dog. J. Serpell. Cambridge, Cambridge University Press.
- Cunningham, M., Downie, M., Eshkar, S., and Kaiser, P. (2001). Loops. Cambridge, MA, MIT Media Lab.
- 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.
- Dautenhahn, K. (1998). "The Art of Designing Socially Intelligent Agents - Science, Fiction, and the Human in the Loop." Special Issue "Socially Intelligent Agents", Applied Artificial Intelligence Journal **12**(7-8).
- Dautenhahn, K. (2000). Socially Intelligent Agents and The Primate Social Brain - Towards a Science of Social Minds. 2000 AAAI Fall Symposium: Socially Intelligent Agents: The Human in the Loop.
- de Vries, H. (1998). "Finding a dominance order most consistent with a linear hierarchy: a new procedure and review." Animal Behaviour **55**: 827-843.
- de Waal, F. B. M., and Harcourt, A.H. (1992). Coalitions and alliances: a history of ethological research. Coalitions and Alliances in Humans and Other Animals. A. H. Harcourt, and de Waal, F. B. M. Oxford, Oxford University press.
- Della Rocca, J. (2001). SIGGRAPH 2001 Summary Report. International Game Developers Association.: [http://www.igda.org/Endeavors/Events/events\\_siggraph2001.htm](http://www.igda.org/Endeavors/Events/events_siggraph2001.htm).
- Dellaert, F., and Beer, R. D. (1994). "Toward an Evolvable Model of Development for Autonomous Agent Synthesis." ALife IV.

- Di Paolo, E. A. (1999). On the Evolutionary and Behavioral Dynamics of Social Coordination: Models and Theoretical Aspects. School of Cognitive and Computing Sciences. Sussex, UK, University of Sussex.
- Dickins, T., and Levy, J. P. (2001). Evolution, development and learning - a nested hierarchy? Connectionist Models of Learning, Development and Evolution: Proceedings of the Sixth Neural Computation and Psychology Workshop. R. French, and Sougne, J. P. London, Springer-Verlag: 263-270.
- Donath, J. (1996). Inhabiting the virtual city: The design of social environments for electronic communities. Media Arts & Sciences. Cambridge, MA, MIT.
- Dowling, J. E. (1992). Neurons and Networks. Cambridge, MA, Belknap Press.
- Downie, M. (2001a). behavior, animation, music: the music and movement of synthetic characters. Media Arts & Sciences. Cambridge, MA, MIT.
- Downie, M. (2001b). music creatures. Proceedings of SIGGRAPH 01: conference abstracts and applications, art gallery.
- Dunbar, R. I. M. (1993). "Coevolution of neocortical size, group size and language in humans." Behavioral and Brain Sciences **16**(4): 681-735.
- Eaton, S. (2001). The Art and Engineering of Variable Character Morphology. Media Arts & Sciences. Cambridge, MA, MIT.
- Ekman, P. (1992). An Argument for Basic Emotions. Basic Emotions. N. Stein, and Oatley, K. Hove, UK, Lawrence Erlbaum: 169-200.
- Engh, A., Esch, L., Smale, L., and Holekamp, K. (2000). "Mechanisms of maternal rank 'inheritance' in the spotted hyaena, *Crocuta crocuta*." Animal Behaviour **60**(323-332).
- Fox, M. W. (1971). Behaviour of Wolves, Dogs and Related Canids. New York, Harper & Row.
- Fragaszy, D., and Visalberghi, E. (1996). Social Learning in Monkeys: Primate "Primacy" Reconsidered. Social Learning in Animals. C. M. Heyes, and Galef, Bennett G., Jr. San Diego, Academic Press.
- Gadanh, S., and Hallam, J. (1998). Emotion-triggered Learning for Autonomous Robots. Simulation of Adaptive Behavior 1998 Workshop on Grounding Emotions in Adaptive Systems.
- Galef, B. G., Jr. (1996). "Food selection: problems in understanding how we choose foods to eat." Neuroscience and Biobehavioral Reviews **20**: 67-73.
- Gaslin, G. (2001). Where the Mona Lisa Meets the Motherboard. Los Angeles Times. 8/17/01. Los Angeles, CA.
- Golani, I., and Moran, G. (1983). A motility-immobility gradient in the behavior of the "inferior" wolf during "ritualized fighting". Advances in the study of mammalian behavior. *Special*

- Publication No. 7. J. F. Eisenberg, and Kleiman, D. G., The American Society of Mammalogists: 65-94.
- Goldberg, D., and Mataric, Maja J. (1997). Interference as a Tool for Designing and Evaluating Multi-Robot Controllers. AAAI-97.
- Goodall, J. (1986). The Chimpanzees of Gombe: Patterns of Behavior. Cambridge, MA, Harvard University Press.
- Gould, J. L. (1982). Ethology. New York, W. W. Norton & Company.
- Grand, S., Cliff, D., and Malhotra, A. (1997). Creatures: artificial life autonomous software agents for home entertainment. Proceedings of the first international conference on Autonomous agents.
- Grenier, M. (2001). Growls and howls hit graphics show. The Wall Street Journal Online: <http://www.zdnet.com/zdn/stories/news/0,4586,2806623,00.html>.
- Gruau, F., and Whitley, D. (1993). "Adding Learning to the Cellular Development of Neural Networks." Evolutionary Computation 1(3): 213-233.
- Harcourt, A. H., and de Waal, F. B. M. (eds.) (1992). Coalitions and Alliances in Humans and other Animals. Oxford, Oxford University Press.
- Hart, W. E., Kammeyer, Thomas E., and Belew, Richard K. (1995). The Role of Development in Genetic Algorithms. Foundations of Genetic Algorithms 3. D. Whitley, and Vose, M. D. California, Morgan Kaufmann.
- Hauser, M. D. (1998). The Evolution of Communication. Cambridge, MA, MIT Press.
- Hayes-Roth, B., Sincoff, E., Brownston, L., Huard, R., and Lent, B. (1995). Directed Improvisation with Animated Puppets. Proceedings of CHI '95.
- Hemelrijk, C. K. (1996). Dominance interactions, spatial dynamics and emergent reciprocity in a virtual world. From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior.
- Hemelrijk, C. K. (1999). "An individual-orientated model of the emergence of despotic and egalitarian societies." Proc. R. Soc. Lond. B 266: 361-369.
- Hemelrijk, C. K. (2000). "Towards the integration of social dominance and spatial structure." Animal Behaviour 59: 1035-1048.
- Heyes, C. M., and Galef, B. G., Jr. (eds.) (1996). Social Learning in Animals. San Diego, Academic Press.
- Hodgins, J. K., and Pollard, N. S. (1997). Adapting Simulated Behaviors For New Characters. Proceedings of the ACM Computer Graphics, SIGGRAPH 97.
- Hogeweg, P. (1988). MIRROR beyond MIRROR, Puddles of LIFE. Artificial life, SFI studies in the sciences of complexity. Redwood City, CA, Addison-Wesley.

- IMAX (1999). Wolves, National Wildlife Federation.
- Isla, D. (2001a). The Virtual Hippocampus: Spatial Common Sense for Synthetic Creatures. Media Arts & Sciences. Cambridge, MA, MIT.
- Isla, D., and Blumberg, B. (2002). AI Game Programming Wisdom. Hingham, MA, Charles River Media.
- Isla, D., Burke, R., Downie, M., and Blumberg, B. (2001b). A Layered Brain Architecture for Synthetic Creatures. Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI).
- Izard, C. E. (1991). The Psychology of Emotions. New York, Plenum Press.
- Jebara, T., Ivanov, Y., Rahimi, A., and Pentland, A. (2000). Tracking Conversational Context for Machine Mediation of Human Discourse. AAAI Fall 2000 Symposium - Socially Intelligent Agents - The Human in the Loop.
- 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 the CHI 99 conference on Human factors in computing systems.
- Johnson-Laird, P., and Oatley, K. (1992). Basic Emotions, Rationality, and Folk Theory. Basic Emotions. N. Stein, and Oatley, K., Lawrence Erlbaum.
- Kawai, M. (1958). "On the system of social ranks in a natural troop of Japanese monkey. I: Basic rank and dependent rank. II: Ranking order as observed among the monkeys on and near the test box." Primates 1: 111-148.
- Kitano, H. (1995). A Model for Hormonal Modulation of Learning. IJCAI.
- Klinghammer, E. (2001). Six day wolf and dog training seminar. Battle Ground, IN, Wolf Park.
- Klinghammer, E., and Goodmann, P. (1985). The Management and Socialization of Captive Wolves (Canis lupus) at Wolf Park. Battle Ground, IN, North American Wildlife Park Foundation, Inc.
- Laland, K. N., Richerson, P. J., and Boyd, R. (1996). Developing a Theory of Animal Social Learning. Social Learning in Animals. C. M. Heyes, and Galef, Bennett G., Jr. San Diego, Academic Press.
- LeDoux, J. (1996). The Emotional Brain. New York, Simon & Schuster.
- Lesser, V. (1995). Proceedings of the First International Conference on Multiagent Systems.
- Lieberman, H., and Vivacqua, A. (2000). Agents to Assist in Finding Help. Proceedings of CHI '00.

- Likert, R. (1932). "A technique for the measurement of attitudes." Archives of Psychology **140**: 1-55.
- Lorenz, K. (1952). King Solomon's Ring. New York, Crowell.
- Lorenz, K., and Leyhausen, P. (1973). Motivation of Human and Animal Behavior: An Ethological View. New York, Van Nostrand Reinhold Co.
- Maes, P. (1989). "How to Do the Right Thing." Connection Science Journal **1**(3).
- Magnenat-Thalmann, N., Kalra, P., and Escher, M. (1998). Communicating with virtual characters. WSCG '98.
- Masuda, T., and Nisbett, R. E. (2001). "Attending Holistically Versus Analytically: Comparing the Context Sensitivity of Japanese and Americans." Journal of Personality and Social Psychology **81**(5).
- Mataric, M. J. (1995). "Issues and Approaches in the Design of Collective Autonomous Agents." Robotics and Autonomous Systems **16**(2-4): 321-331.
- Matsuzawa, T. (1994). Field Experiments on Use of Stone Tools by Chimpanzees in the Wild. Chimpanzee Cultures. R. W. Wrangham, McGrew, W. C., de Waal, Frans B. M., and Heltne, Paul G. Cambridge, Harvard University Press.
- Maynard Smith, J. (1964). "Group selection and kin selection." Nature **201**(4924): 1145-1147.
- McNeill, D. (1998). The Face. New York, Little Brown & Company.
- Mech, L. D. (1970). The Wolf. Minneapolis, MN, University of Minnesota Press.
- Mech, L. D. (1988). The arctic wolf: living with the pack. Stillwater, MN, Voyageur Press.
- Mech, L. D. (1999). "Alpha status, dominance, and division of labor in wolf packs." Canadian Journal of Zoology **77**: 1196-1203.
- 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.
- Mehrabian, A., and Russell, J. (1974). An Approach to Environmental Psychology. Cambridge, MA, MIT Press.
- Michaud, F., Robichaud, E., and Audet, J. (2001). Using Motives and Artificial Emotions for Prolonged Activity of a Group of Autonomous Robots. North Falmouth, MA, AAAI Fall Symposium: 85-90.
- Minsky, M. (1986). The Society of Mind. New York, Simon and Schuster.
- Minsky, M., ed. (1968). Semantic information processing. Cambridge, MA, MIT Press.

- Mitani, J. C., Merriwether, D.A. and Zhang, C. (2000). "Male affiliation, cooperation and kinship in wild chimpanzees." Animal Behaviour **59**: 885-893.
- Morss, C. (2000). Seminar on wolf behavior, Wolf Hollow Ipswich.
- Murie, A. (1944). The Wolves of Mount McKinley. Fauna of the National Parks Series. Washington, DC, U.S. National Park Service. **5**.
- Nass, C., Fogg, B. J., and Moon, Y. (1996). "Can computers be teammates?" International Journal of Human-Computer Studies **45**(6): 669-678.
- Nass, C., Steuer, J. S., and Tauber, E. (1994). Computers are social actors. CHI '94 Proceedings.
- Noble, J. (1998). The Evolution of Animal Communication Systems: Questions of Function Examined through Simulation. Sussex, UK, University of Sussex.
- Noble, J., Tuci, E., Todd, P. M. (1999). Social learning and information sharing: an evolutionary simulation model of foraging in Norway rats. Proc. of the Fifth European Conference on Artificial Life (ECAL'99).
- OED (2002). Oxford English Dictionary: <http://dictionary.oed.com/cgi/entry/00229735>.
- Olsen, G. (2001). SIGGRAPH 2001 - is experimentation really dangerous?, Interaction By Design, 8/17/01: <http://www.interactionbydesign.com/thoughts/thumbnails/00000045.html>.
- Packard, J. (in press). Social Behavior of Wolves. The Ecology and Behavior of the Wolf. L. D. a. B. Mech, L. Chicago, The University of Chicago Press.
- Packer, C. (1977). "Reciprocal altruism in *Papio anubis*." Nature **265**: 441-443.
- Pepperberg, I. M. (1999). The Alex studies: cognitive and communicative abilities of grey parrots. Cambridge, MA, Harvard University Press.
- Perlin, K., and Goldberg, A. (1996). Improv: A System for Scripting Interactive Actors in Virtual Worlds. Proceedings of the ACM Computer Graphics, SIGGRAPH 96.
- Picard, R. (1997). Affective Computing. Cambridge, MA, MIT Press.
- Plotkin, H. (1994). Darwin Machines and the Nature of Knowledge. Cambridge, MA, Harvard University Press.
- Plutchik, R. (1991). The Emotions. Lanham, MD, University Press of America.
- Reeves, B., and Nass, C. (1996). The Media Equation. Cambridge, Cambridge University Press.
- Reilly, W. S. N. (1996). Believable Social and Emotional Agents. School of Computer Science. Pittsburgh, PA, Carnegie Mellon University.
- Resnick, M. (1994). Turtles, Termites, and Traffic Jams. Cambridge, MA, MIT Press.

- Reynolds, C. (1987). Flocks, Herds and Schools: A Distributed Behavioral Model. Proceedings of the ACM Computer Graphics, 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. (1980). "A circumplex model of affect." Journal of Personality and Social Psychology **39**: 1161-1178.
- 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.
- Ryokai, K., Vaucelle, C., and Cassell, J. (in press). Literacy Learning by Storytelling with a Virtual Peer. Proceedings of Computer Support for Collaborative Learning.
- 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.
- Serpell, J. (1995). The Domestic Dog. Cambridge, Cambridge University Press.
- Shannon, C. E. (1948). "A Mathematical Theory of Communication." The Bell System Technical Journal **27**(July, October, 1948): 379-423, 623-656.
- Shardanand, U., and Maes, P. (1995). Social Information Filtering: Algorithms for Automating 'Word of Mouth'. Proceedings of CHI '95.
- Siegel, D. J. (1999). The Developing Mind. New York, The Guilford Press.
- Smith, C. (1989). "Dimensions of appraisal and physiological response in emotion." Journal of Personality and Social Psychology **56**: 339-353.
- Steels, L., and Brooks, R. A., editors (1995). The Artificial Life Route to Artificial Intelligence: Building Embodied Situated Agents. Hillsdale, NJ, Lawrence Erlbaum Associates, Inc.
- Stone, P. (2000). Layered Learning in Multi-Agent Systems: A Winning Approach to Robotic Soccer. Cambridge, MA, MIT Press.
- Strassman, S. (1994). Semi-autonomous animated actors. Proceedings of the 12th National Conference on Artificial Intelligence.
- Stroud, M. (2001). Stretching it out at SIGGRAPH, Wired News, 8/11/01: <http://www.wired.com/news/culture/0,1284,46005,00.html>.

- Sutton, R., and Barto, A. (1998). Reinforcement Learning: An Introduction. Cambridge, MA, MIT Press.
- TCSGlobalNews (1996). <http://www.telecoms-mag.com/marketing/articles/aug96/globe.html>. **2000**.
- teBoekhorst, I., and Hogeweg, P. (1994). "Self-structuring in artificial 'CHIMPS' offers new hypothesis for male grouping in chimpanzees." Behaviour **130**: 229-252.
- Terkel, J. (1996). Cultural Transmission of Feeding Behavior in the Black Rat (*Rattus rattus*). Social Learning in Animals. C. M. Heyes, and Galef, Bennett G., Jr. San Diego, Academic Press.
- 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.
- Thomas, F., and Johnson, O. (1981). Disney Animation: The Illusion of Life. New York, Abbeville Press.
- Thorndike, E. (1911). Animal Intelligence. Darien, CT, Hafner.
- Tomlinson, B. (2000a). "Dead Technology." Style **22**(2): 316-335.
- Tomlinson, B., Blumberg, B., and Nain, D. (2000b). "Expressive Autonomous Cinematography for Interactive Virtual Environments." Proceedings of 4th Int'l. Conf. on Autonomous Agents.
- Tomlinson, B., Downie, M., Berlin, M., Gray, J., Wong, A., Burke, R., Isla, D., Ivanov, Y., Johnson, M. P., Lyons, D., Cochran, J., Yong, B., Stiehl, D., Soetjipto, R., Zaharopol, D., and Blumberg, B. (2001). AlphaWolf. Proceedings of SIGGRAPH 2001: conference abstracts and applications.
- Trivers, R. L. (1971). "The evolution of reciprocal altruism." Quarterly Review of Biology **46**: 35-57.
- Tu, X., and Terzopoulos, D (1994). Artificial Fishes: Physics, Locomotion, Perception, Behavior. Proceedings of the ACM Computer Graphics, SIGGRAPH 94.
- Vaario, J. (1991). Developing neural networks through growth. Proceedings of the 5th Annual Conference of JSAI (Japanese Society for Artificial Intelligence).
- Vaughan, R. T., Stoy, K., Sukhatme, G.S., Mataric, M.J. (2000). "Go ahead, make my day: robot conflict resolution by aggressive competition." Proc. Int. Conf. Simulation of Adaptive Behaviour.
- Velasquez, J. (1998). "When Robots Weep: Emotional Memories and Decision-Making." Proceedings of the Fifteenth National Conference on Artificial Intelligence.
- Visalberghi, E., and Addessi, E. (2000). "Seeing group members eating a familiar food enhances the acceptance of novel foods in capuchin monkeys." Animal Behaviour **60**: 69-76.



- Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., and Thelen, E. (2001). "Autonomous Mental Development by Robots and Animals." Science **291**(5504): 599-900.
- Wilson, E. O. (1980). Sociobiology. Cambridge, MA, Belknap Press of Harvard University.
- Wrangham, R., and Peterson, D. (1996). Demonic Males. Boston, Houghton Mifflin.
- Wrangham, R. W., McGrew, W.C., de Waal, F.B.M., and Heltne, P.G., eds. (1994). Chimpanzee Cultures. Cambridge, MA, Harvard University Press.
- www.wolf.org (2000). Frequently Asked Questions, International Wolf Center. **2001**.
- Yoon, S.-Y. (2000a). Affective Synthetic Characters. Media Laboratory. Cambridge, MA, MIT.
- Yoon, S.-Y., Blumberg, B., and Schneider, G. (2000b). "Motivation Driven Learning for Interactive Synthetic Characters." Proceedings of 4th Int'l. Conf. on Autonomous Agents.
- Zabel, C., Glickman, S., Frank, L., Woodmansee, K. and Keppel, G. (1992). Coalition formation in a colony of prepubertal spotted hyenas. Coalitions and Alliances in Humans and Other Animals. A. H. Harcourt, and de Waal, F. B. M. Oxford, Oxford University press.
- Zentall, T. (1996). An Analysis of Imitative Learning in Animals. Social Learning in Animals. C. M. Heyes, and Galef, Bennett G., Jr. San Diego, Academic Press.
- Zimen, E. (1981). The Wolf. New York, Delacorte Press.

## Appendix: Experimental Protocol

This Appendix contains most of the information about the experimental method of the AlphaWolf user studies.

### *Experimenter's Script*

#### PRE-EXPERIMENT PREPARATION

*[Make sure you have the master list of what experiments they'll be doing.*

*Set up batch files to launch random-ordered experiments.*

*Check creation dates on batch files to confirm that it worked.*

*Verify sound in FOR BASS and out – CD on both.*

*Remove command and relationship cards.*

*Put ID # and Run # on all forms.*

*Mark date and time on master list.*

*Verify room is clean and curtains are up.]*

#### INTRODUCTION

“Thanks for helping out with this experiment about social relationships in virtual wolves. First, I'm going to tell you about what you're going to be doing in the study.”

“There are several parts to this experiment, each of which will last approximately 5 minutes. First, you will be asked to watch a short video of real wolves in the wild. Then you will watch or participate in three runs of computer graphical wolves. After each run, you will be asked to fill out an identical short questionnaire. At the end of the experiment, there will be one additional questionnaire, comparing the runs. The entire experiment will last approximately 45 minutes. You will be given a \$10 gift certificate to Toscanini's as compensation for your participation.”

“You are welcome to stop at any time. All information collected during your participation will be destroyed and your payment will be prorated based on the time you have already spent. Do you have any questions?” *[Discuss any questions the subject has.]*

“Now I need you to sign this consent form.” *[Hand subject Consent Form.]* “Please read the consent form now, and if you agree, please sign at the bottom.” *[Give subject time to read and sign Consent Form.] [Collect Consent Form and put in folder.]*

*[Hand subject questionnaire.]* “At the end of each run of virtual wolves, you'll be asked to fill out this questionnaire. Most of the questions have a seven point scale for answers. When it's time to fill it out, please circle the number that corresponds to your answer. Please take a moment now to look it over, so that you know what to look for in both the real and virtual wolves.” *[Give subject time to read the questionnaire.] [Hand subject notepad.]* “Here's a notepad, in case you want to write anything down at any point. Do you have any questions? Would you like a glass of water or anything?” *[Discuss any questions the subject has. Get water, if requested.]*

#### REAL WOLVES

“If you're ready, I'll start the video of real wolves in the wild. This is an excerpt from a National Geographic special. It will play for approximately three minutes. I'll come back in when it finishes.” *[Play video and leave room.]*

## **VIRTUAL WOLVES**

“Now we’ll do the virtual wolves. There will be three different runs. Each run has three wolf pups in it, one gray, one black, and one white.”

**VIDEO:** “This run is a short video of some virtual wolves on this computer screen. The wolves will play for approximately four minutes. *[Hand subject questionnaire.]* Here’s a questionnaire for this section. You’ll have as much time as you need at the end to fill out the questionnaire. You’ll know the wolves have finished loading because you’ll start to hear the sound of wind. When the four minutes are up, the run will stop automatically. When it finishes, please fill out the questionnaire. I’ll be outside the door; please let me know when you’re done with the questionnaire, or come get me if you have any questions. Do you have any questions now?” *[Discuss any questions the subject has.]* “If you’re ready to begin, I’ll start the run now.” *[Play virtual wolves, and leave as soon as the wind comes on.] [When subject is finished, collect questionnaire.]*

**BUTTONS:** “Now you’ll play a virtual wolf game. You’ll be playing the role of the gray pup. The goal of the game is to get your pup to form relationships with its littermates. *[Give subject controls sheet.]* Your controls are as follows: clicking with the mouse makes your pup run to where you clicked. Clicking on another wolf causes your pup to interact with that wolf. Clicking on the picture of the wolf at the bottom of the screen will make your pup go find that wolf and interact with it. The icons will light up when your pup’s interacting with that individual.”

“You can control what action your pup takes by hitting these buttons on the keyboard. Pressing the key labeled “GROWL” makes your pup try to dominate another wolf. Pressing the key labeled “WHINE” makes your pup submit to that wolf.”

“You’ll have four minutes to form your relationships with the other pups. *[Give card.]* This card shows you what relationship you should form with each of your siblings. Your goal is to try to *[read card]* {dominate/submit to} the white pup, and {dominate/submit to} the black pup.

*[Hand subject questionnaire.]* “Here’s a questionnaire for this section. You’ll have as much time as you need at the end to fill out the questionnaire.”

“It’ll take about 60 seconds to load the wolves. You’ll know the wolves have finished loading because you’ll start to hear the sound of wind. Your pup will start off asleep. You can wake it up by clicking it with the mouse or hitting the marked keys. When the four minutes are up, the run will stop automatically. When it finishes, please fill out the questionnaire. I’ll be outside the door; please let me know when you’re done with the questionnaire, or come get me if you have any questions. Do you have any questions now?” *[Discuss any questions the subject has.]* “If you’re ready to begin, I’ll start the run now.” *[Play virtual wolves, and leave as soon as the wind comes on.] [When subject is finished, collect questionnaire.]*

**MICROPHONE:** “Now you’ll play a virtual wolf game. You’ll be playing the role of the gray pup. The goal of the game is to get your pup to form relationships with its littermates. *[Give subject controls sheet.]* Your controls are as follows: clicking with the mouse makes your pup run to where you clicked. Clicking on another wolf causes your pup to interact with that wolf. Clicking on the picture of the wolf at the bottom of the screen will make your pup go find that wolf and interact with it. The icons will light up when your pup’s interacting with that individual.”

“Growling into this microphone makes your pup try to dominate another wolf. Whining into it makes your pup submit to that wolf. In general you should hold it about this far from your mouth. Growls sound like this *[demonstrate]*, and whines sound like this: *[demonstrate]*. If you have trouble growling, barking works well for dominating, too. Human speech and laughter gets picked up as whining, so try not to laugh into the microphone or you’ll get a very submissive pup.”

“You’ll have four minutes to form your relationships with the other pups. *[Give card.]* This card shows you what relationship you should form with each of your siblings. Your goal is to try to *[read card]* {dominate/submit to} the white pup, and {dominate/submit to} the black pup.

*[Hand subject questionnaire.]* “Here’s a questionnaire for this section. You’ll have as much time as you need at the end to fill out the questionnaire.”

“It’ll take about 60 seconds to load the wolves. You’ll know the wolves have finished loading because you’ll start to hear the sound of wind. Your pup will start off asleep. You can wake it up by clicking it with the mouse or hitting the marked keys. When the four minutes are up, the run will stop automatically. When it finishes, please fill out the questionnaire. I’ll be outside the door; please let me know when you’re done with the questionnaire, or come get me if you have any questions. Do you have any questions now?” *[Discuss any questions the subject has.]* “If you’re ready to begin, I’ll start the run now.” *[Play virtual wolves, and leave as soon as the wind comes on.] [When subject is finished, collect questionnaire.]*

## **ENDING**

Now I’ll ask you to fill out one final questionnaire, comparing the various runs.” *[Give subject the correct one of the two ending questionnaires.]* “I’ll give you a few minutes to fill it out. Please let me know when you’re done.” *[When subject is finished, collect questionnaire.]*

“That’s the end of the experiment! Thanks for participating. Here’s your Toscanini’s gift certificate. I’ll now read you a Debriefing Statement to explain the purpose of the study.” *[Read Debriefing Statement.]* “Do you have any final questions?” *[Discuss any questions the subject has.]* “Thanks again for your time.”

## ***Consent Form***

User ID: \_\_\_\_\_

This form is designed to provide you with information about this study. Your participation in the following experiment is completely voluntary. You are free to withdraw this consent at any time, for any reason, and to request that any data collected be destroyed. If at any time you feel uncomfortable, or unsure that you wish your results to be part of the experiment, you may discontinue your participation with no repercussions. All information collected during your participation will be destroyed and your payment will be prorated based on the time you have already spent.

There are several parts to this experiment, each of which will last approximately 5 minutes. First, you will be asked to watch a short video of real wolves in the wild. Then you will watch or participate in three runs of computer graphical wolves. After each run, you will be asked to fill out an identical short questionnaire. At the end of the experiment, there will be one additional questionnaire, comparing the runs. The entire experiment will last approximately 45 minutes. You will be given a \$10 gift certificate to Toscanini's as compensation for your participation.

Any responses that are collected during the experiment will be completely anonymous. From this point forward, only the ID number that appears on the upper right corner of this packet will be used to refer to you. If you have any questions at any point during the experiment, the experimenter will gladly answer them.

Please read the following and sign on the lines below:

"I, the undersigned, have read and understood the explanations of the following research project and voluntarily consent to my participation in it. I understand that my responses will remain confidential and that I may terminate my participation at any time.

In the unlikely event of physical injury resulting from participation in this research, I understand that medical treatment will be available from the MIT Medical Department, including first aid emergency treatment and follow-up care as needed, and that my insurance carrier may be billed for the cost of such treatment. However, no compensation can be provided for medical care apart from the foregoing. I further understand that making such medical treatment available; or providing it, does not imply that such injury is the Investigator's fault. I also understand that by my participation in this study I am not waiving any of my legal rights.

I understand that I may also contact the Chairman of the Committee on the Use of Humans of Experimental Subjects, MIT 253-6787, if I feel I have been treated unfairly as a subject.

Signature: \_\_\_\_\_

Name (Printed): \_\_\_\_\_

Date: \_\_\_\_\_

Location: NE18-5FL

## Questionnaire

Questionnaire #: \_\_\_\_\_

User ID: \_\_\_\_\_

1. Gender [M] [F]

2. Age \_\_\_\_\_

**Please circle the number that corresponds to your answer. Except where specified, all questions pertain to the run of virtual wolves that you have just watched.**

3. At the end of this run, who was the more dominant wolf pup – white or gray?

**White is dominant**      1      2      3      4      5      6      7      **Gray is dominant**

4. At the end of this run, who was the more dominant wolf pup – white or black?

**White is dominant**      1      2      3      4      5      6      7      **Black is dominant**

5. At the end of this run, who was the more dominant wolf pup – gray or black?

**Gray is dominant**      1      2      3      4      5      6      7      **Black is dominant**

6. How clear were the social dominance relationships between the virtual wolves?

**Not clear at all**      1      2      3      4      5      6      7      **Very clear**

7. How much did the relationships between each pair of virtual wolves change over the course of the run?

**Not at all**      1      2      3      4      5      6      7      **Very much**

8. How much control did you feel you had over the relationships in this run?

**None at all**      1      2      3      4      5      6      7      **Very much**

9. How much control did you feel you had over the behavior of the white pup (either directly or indirectly)?

**None at all**      1      2      3      4      5      6      7      **Very much**

10. How much control did you feel you had over the behavior of the gray pup (either directly or indirectly)?

**None at all**      1      2      3      4      5      6      7      **Very much**

11. How much control did you feel you had over the behavior of the black pup (either directly or indirectly)?

**None at all**      1      2      3      4      5      6      7      **Very much**

12. How similar was the overall social behavior of the virtual wolves to the behavior of real wolves (for example, the ones that you saw in the video)?

**Not similar at all**      1      2      3      4      5      6      7      **Very similar**

*(please turn over)*

13. How similar was the submissive behavior of the virtual wolves to the behavior of real wolves (for example, the ones that you saw in the video)?

**Not similar at all**      1      2      3      4      5      6      7      **Very similar**

14. How similar was the dominance behavior of the virtual wolves to the behavior of real wolves (for example, the ones that you saw in the video)?

**Not similar at all**      1      2      3      4      5      6      7      **Very similar**

15. How much did you enjoy this run of the virtual wolves?

**Not at all**      1      2      3      4      5      6      7      **Very much**

16. How much did you like this pack of virtual wolves?

**Not at all**      1      2      3      4      5      6      7      **Very much**

17. How much did you like the white pup?

**Not at all**      1      2      3      4      5      6      7      **Very much**

18. How much did you like the gray pup?

**Not at all**      1      2      3      4      5      6      7      **Very much**

19. How much did you like the black pup?

**Not at all**      1      2      3      4      5      6      7      **Very much**

20. How much did you identify with the white pup?

**Not at all**      1      2      3      4      5      6      7      **Very much**

21. How much did you identify with the gray pup?

**Not at all**      1      2      3      4      5      6      7      **Very much**

22. How much did you identify with the black pup?

**Not at all**      1      2      3      4      5      6      7      **Very much**

23. On the rest of this page, please describe any additional thoughts you had on the virtual wolves.

## *Final Questionnaire – Microphone/Button/Video*

User ID: \_\_\_\_\_

**Please circle the number that corresponds to your answer. In all of the questions, the run that you watched but did not interact with will be called the “Video Run”, the run where you pressed buttons to direct the wolf’s actions will be called “Button Run”, and the run where you used the microphone will be called “Microphone Run”.**

1. Which of the runs did you enjoy more?

**Video Run**            1            2            3            4            5            6            7            **Button Run**

2. Which of the runs did you enjoy more?

**Microphone Run**            1            2            3            4            5            6            7            **Button Run**

3. Which of the runs did you enjoy more?

**Microphone Run**            1            2            3            4            5            6            7            **Video Run**

4. In which of the runs did you feel you had more control over the behavior of the wolves?

**Video Run**            1            2            3            4            5            6            7            **Button Run**

5. In which of the runs did you feel you had more control over the behavior of the wolves?

**Microphone Run**            1            2            3            4            5            6            7            **Button Run**

6. In which of the runs did you feel you had more control over the behavior of the wolves?

**Microphone Run**            1            2            3            4            5            6            7            **Video Run**

7. In which of the runs did the behavior of the virtual wolves more closely resemble the behavior of real wolves?

**Video Run**            1            2            3            4            5            6            7            **Button Run**

8. In which of the runs did the behavior of the virtual wolves more closely resemble the behavior of real wolves?

**Microphone Run**            1            2            3            4            5            6            7            **Button Run**

9. In which of the runs did the behavior of the virtual wolves more closely resemble the behavior of real wolves?

**Microphone Run**            1            2            3            4            5            6            7            **Video Run**

10. In which of the runs did you feel more comfortable with the interface or lack thereof?

**Video Run**            1            2            3            4            5            6            7            **Button Run**

11. In which of the runs did you feel more comfortable with the interface or lack thereof?

**Microphone Run**            1            2            3            4            5            6            7            **Button Run**

*(please turn over)*



12. In which of the runs did you feel more comfortable with the interface or lack thereof?  
**Microphone Run**      1      2      3      4      5      6      7      **Video Run**
13. In which of the runs were the social dominance relationships clearer to you?  
**Video Run**      1      2      3      4      5      6      7      **Button Run**
14. In which of the runs were the social dominance relationships clearer to you?  
**Microphone Run**      1      2      3      4      5      6      7      **Button Run**
15. In which of the runs were the social dominance relationships clearer to you?  
**Microphone Run**      1      2      3      4      5      6      7      **Video Run**
16. In which of the runs did you feel more immersed in the experience?  
**Video Run**      1      2      3      4      5      6      7      **Button Run**
17. In which of the runs did you feel more immersed in the experience?  
**Microphone Run**      1      2      3      4      5      6      7      **Button Run**
18. In which of the runs did you feel more immersed in the experience?  
**Microphone Run**      1      2      3      4      5      6      7      **Video Run**
19. In which of the runs did you like the wolves more?  
**Video Run**      1      2      3      4      5      6      7      **Button Run**
20. In which of the runs did you like the wolves more?  
**Microphone Run**      1      2      3      4      5      6      7      **Button Run**
21. In which of the runs did you like the wolves more?  
**Microphone Run**      1      2      3      4      5      6      7      **Video Run**
22. On the rest of this page, please describe any additional thoughts you had on any aspect of this experiment.

## *Final Questionnaire – Button 1/Button 2/Video*

User ID: \_\_\_\_\_

**Please circle the number that corresponds to your answer. In all of the questions, the run that you watched but did not interact with will be called the “Video Run”, the first run where you pressed buttons to direct the wolf’s actions will be called “Button Run 1”, and the second run where you pressed buttons to direct the wolf’s actions will be called “Button Run 2”.**

1. Which of the runs did you enjoy more?

**Video Run**            1        2        3        4        5        6        7        **Button Run 1**

2. Which of the runs did you enjoy more?

**Button Run 2**            1        2        3        4        5        6        7        **Button Run 1**

3. Which of the runs did you enjoy more?

**Button Run 2**            1        2        3        4        5        6        7        **Video Run**

4. In which of the runs did you feel you had more control over the behavior of the wolves?

**Video Run**            1        2        3        4        5        6        7        **Button Run 1**

5. In which of the runs did you feel you had more control over the behavior of the wolves?

**Button Run 2**            1        2        3        4        5        6        7        **Button Run 1**

6. In which of the runs did you feel you had more control over the behavior of the wolves?

**Button Run 2**            1        2        3        4        5        6        7        **Video Run**

7. In which of the runs did the behavior of the virtual wolves more closely resemble the behavior of real wolves?

**Video Run**            1        2        3        4        5        6        7        **Button Run 1**

8. In which of the runs did the behavior of the virtual wolves more closely resemble the behavior of real wolves?

**Button Run 2**            1        2        3        4        5        6        7        **Button Run 1**

9. In which of the runs did the behavior of the virtual wolves more closely resemble the behavior of real wolves?

**Button Run 2**            1        2        3        4        5        6        7        **Video Run**

10. In which of the runs did you feel more comfortable with the interface or lack thereof?

**Video Run**            1        2        3        4        5        6        7        **Button Run 1**

11. In which of the runs did you feel more comfortable with the interface or lack thereof?

**Button Run 2**            1        2        3        4        5        6        7        **Button Run 1**

*(please turn over)*

12. In which of the runs did you feel more comfortable with the interface or lack thereof?  
**Button Run 2**            1            2            3            4            5            6            7            **Video Run**
13. In which of the runs were the social dominance relationships clearer to you?  
**Video Run**            1            2            3            4            5            6            7            **Button Run 1**
14. In which of the runs were the social dominance relationships clearer to you?  
**Button Run 2**            1            2            3            4            5            6            7            **Button Run 1**
15. In which of the runs were the social dominance relationships clearer to you?  
**Button Run 2**            1            2            3            4            5            6            7            **Video Run**
16. In which of the runs did you feel more immersed in the experience?  
**Video Run**            1            2            3            4            5            6            7            **Button Run 1**
17. In which of the runs did you feel more immersed in the experience?  
**Button Run 2**            1            2            3            4            5            6            7            **Button Run 1**
18. In which of the runs did you feel more immersed in the experience?  
**Button Run 2**            1            2            3            4            5            6            7            **Video Run**
19. In which of the runs did you like the wolves more?  
**Video Run**            1            2            3            4            5            6            7            **Button Run 1**
20. In which of the runs did you like the wolves more?  
**Button Run 2**            1            2            3            4            5            6            7            **Button Run 1**
21. In which of the runs did you like the wolves more?  
**Button Run 2**            1            2            3            4            5            6            7            **Video Run**
22. On the rest of this page, please describe any additional thoughts you had on any aspect of this experiment.

## ***Debriefing Statement***

The experiment you just participated in was designed to test several mechanisms of social relationship formation among computational entities. Various subjects viewed and interacted with virtual wolves with different mechanisms of social relationship formation. We are trying to discover the extent to which these mechanisms capture what people and wolves think of as “social relationships”. We will be analyzing your responses to the questionnaires to see how the various mechanisms compare to each other.

The VIDEO pack of virtual wolves that you watched were the result of somebody else playing the wolf game. We’ll be using one of the games that you just played in someone else’s experiment in a few days. The game will be completely anonymous – no identifying information will be attached to the game.

If at any time, now or later, you experience any ill effects (either mental or physical) as a result of your participation in this experiment, please do not hesitate to tell the experimenter, or call 617-452-5611 and ask for Bill.

Your help has been greatly appreciated, and will aid our understanding of how to make computational entities that can form social relationships with each other and with people.

I’d just like to add one more thing before you go. It’s really important that you not discuss this experiment with anyone else, because if someone in the future does the experiment knowing what we are trying to study they can throw the results off. Thanks!